



Trigger & DAQ R&D for high rate experiments



Wesley H. Smith

U. Wisconsin – Madison

*CMS Trigger Coordinator &
Upgrade R&D Peer Review Chair*

Workshop on Detector R&D

Fermilab

October 8, 2010

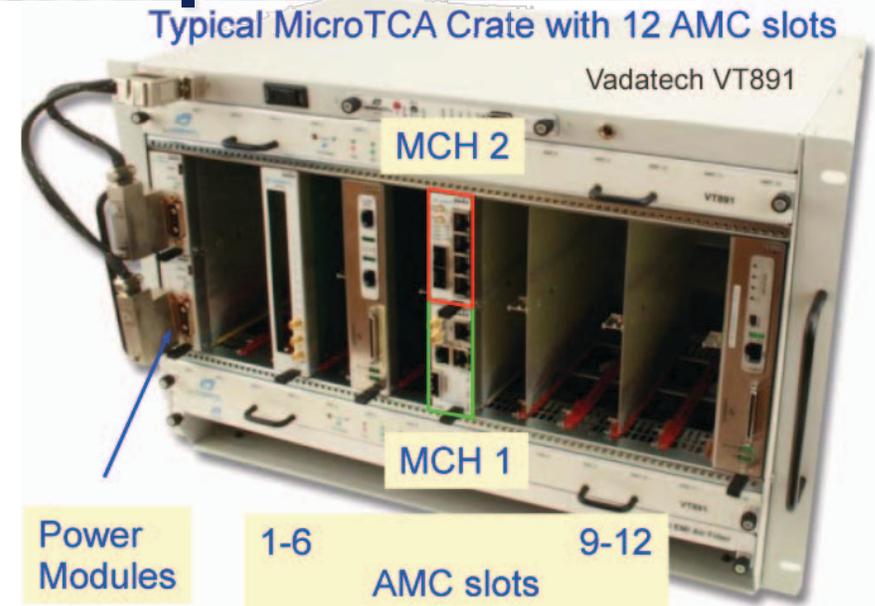
Outline:

- Tools for High Rate Experiments: μ TCA, FPGAs, Transceivers
- Mu2e Trigger & DAQ
- ATLAS & CMS Trigger & DAQ
- LHC evolution & challenges
- Upgrades for ATLAS & CMS Trigger & DAQ



HEP tools for high rate experiments: μ TCA

- **Advanced Telecommunications Computing Architecture ATCA**
- **μ TCA Derived from AMC std.**
 - **Advanced Mezzanine Card**
 - **Up to 12 AMC slots**
 - *Processing modules*
 - **1 or 2 MCH slots**
 - *Controller Modules*
- **6 standard 10Gb/s point-to-point links from each slot to hub slots (more available)**
- **Redundant power, controls, clocks**
- **Each AMC can have in principle (20) 10 Gb/sec ports**
- **Backplane customization is routine & inexpensive**



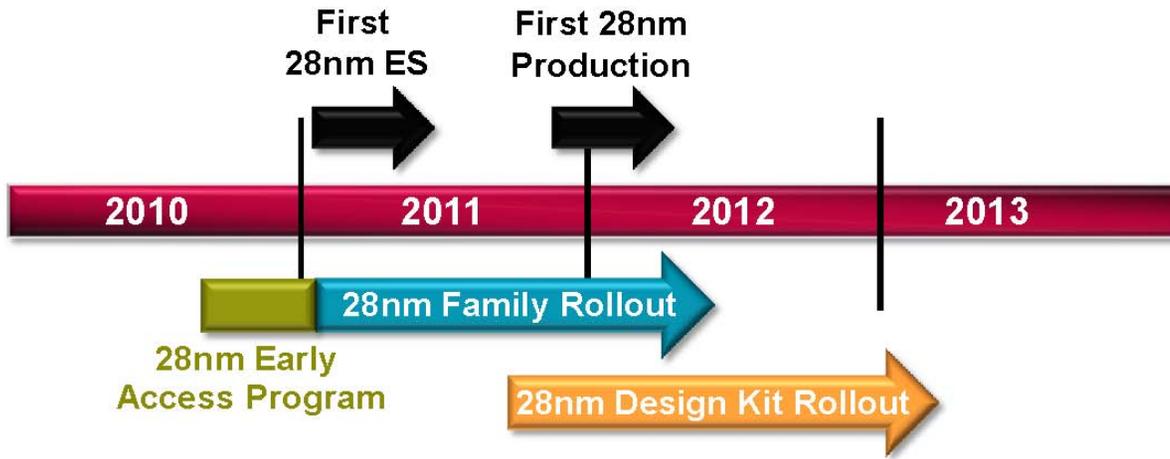
Single Module (shown): 75 x 180 mm
Double Module: 150 x 180mm





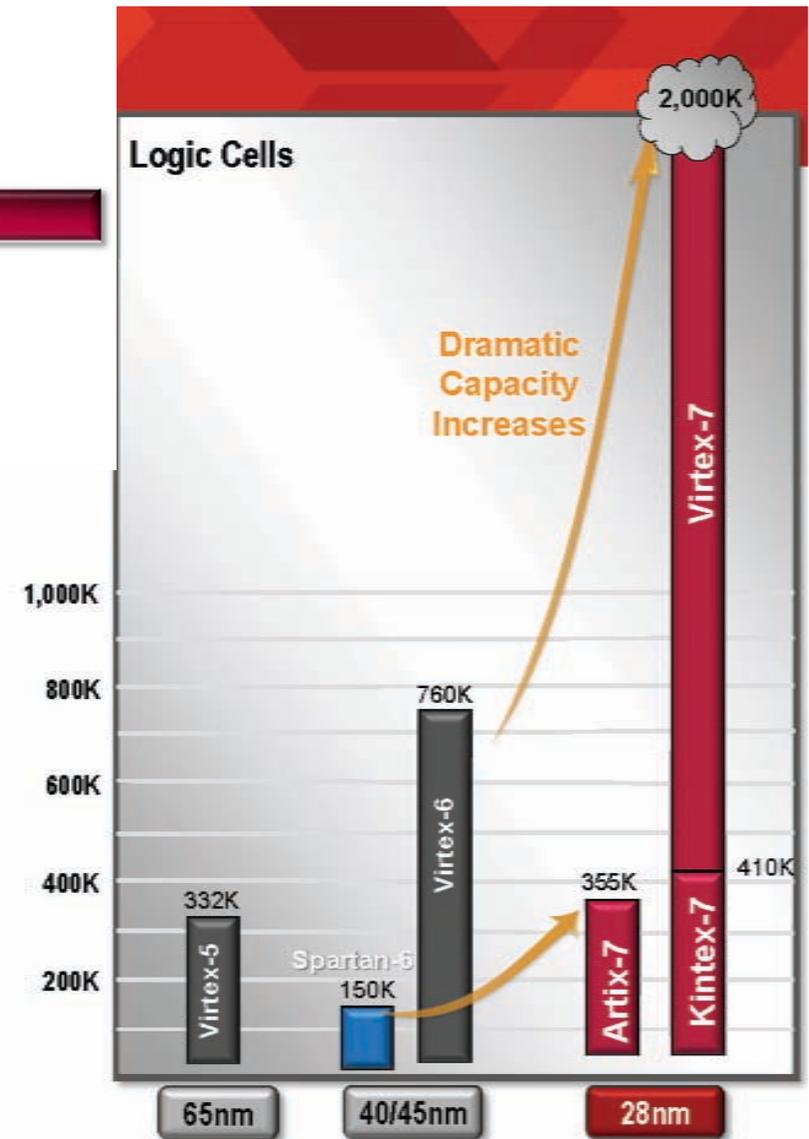
Workshop on Detector R&D

FPGAs: Logic Cells



Next generation 28 nm:
 > 2x capacity gains over 40nm devices

Family	Capacity Range
Artix-7	20K – 355K LCs
Kintex-7	30K – 410K LCs
Virtex-7	285K – 2,000K LCs





FPGAs: Transceivers



Challenge:

- Increase device BW
- No increase in total device power
- XCVR gains from scaling: negligible

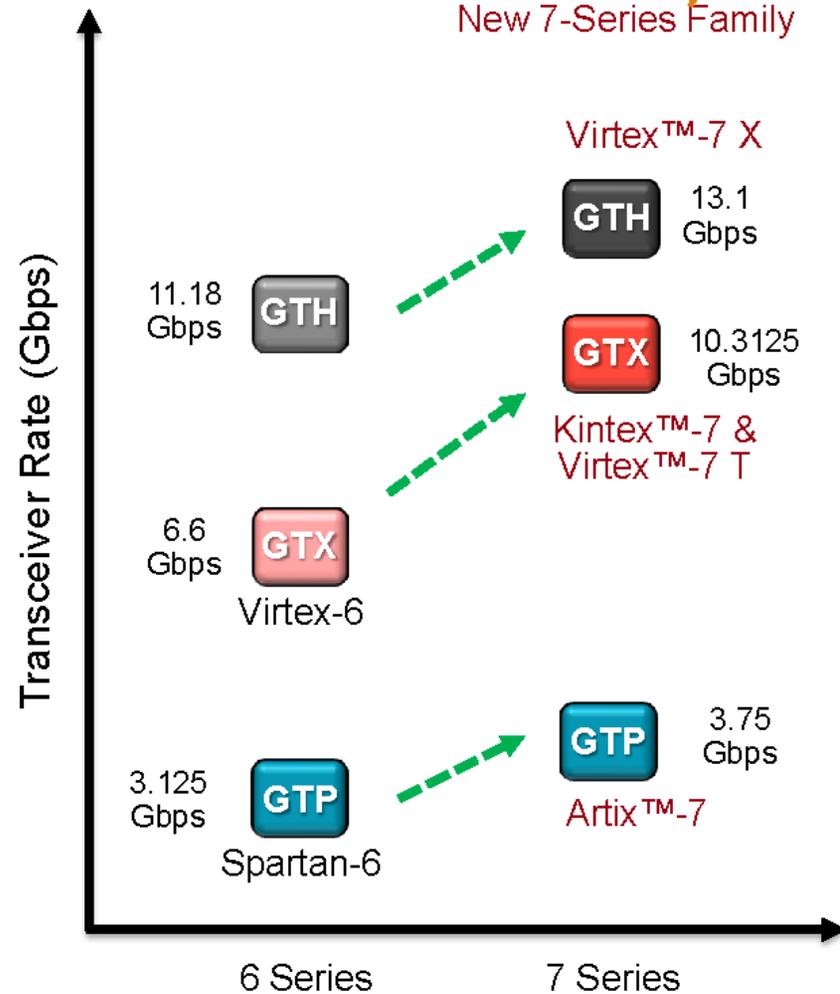
Solution:

- Careful circuit design throughout XCVR
- Increased Gbps / XCVR
- More XCVR / Device
- Low power mode for short channels
- Lanes share a PLL vs PLL per lane

Result:

- 60% Increased max device BW
- Device XCVR power unchanged

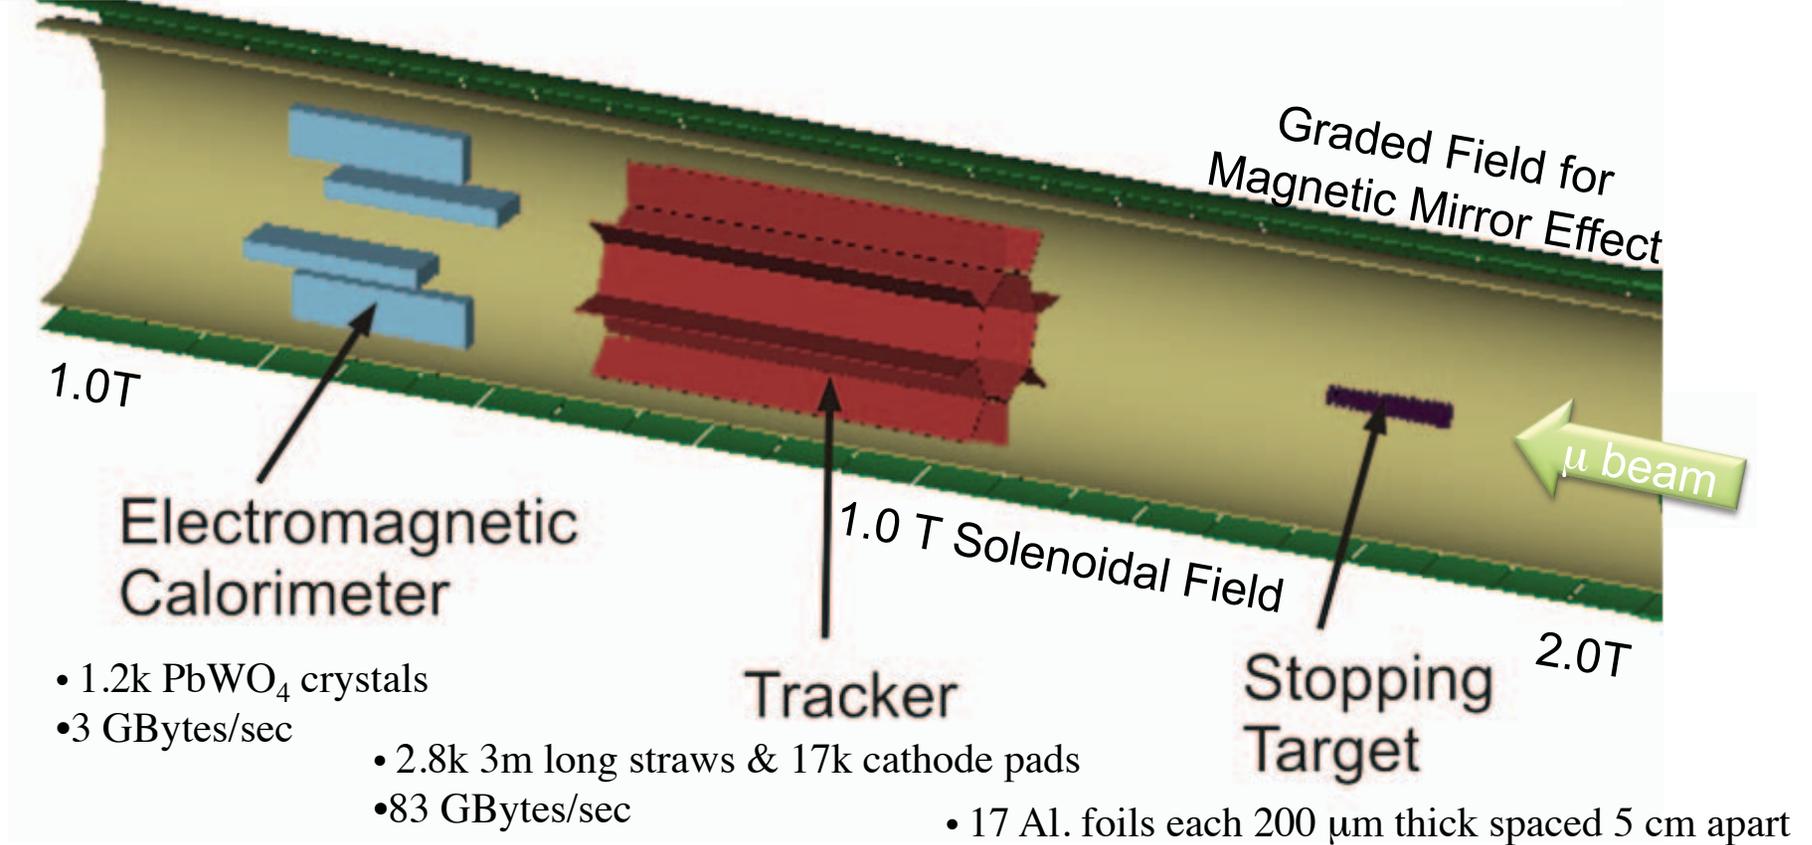
	GTP	GTX	GTH	GT28
Max Rate (Gbps)	3.75	10.3125	13.1	28
Relative Power (Per GT)	.35x	.7x	1x	-
Max GTs per Device	4	56	72	-





Workshop on Detector R&D

Mu2e Detector

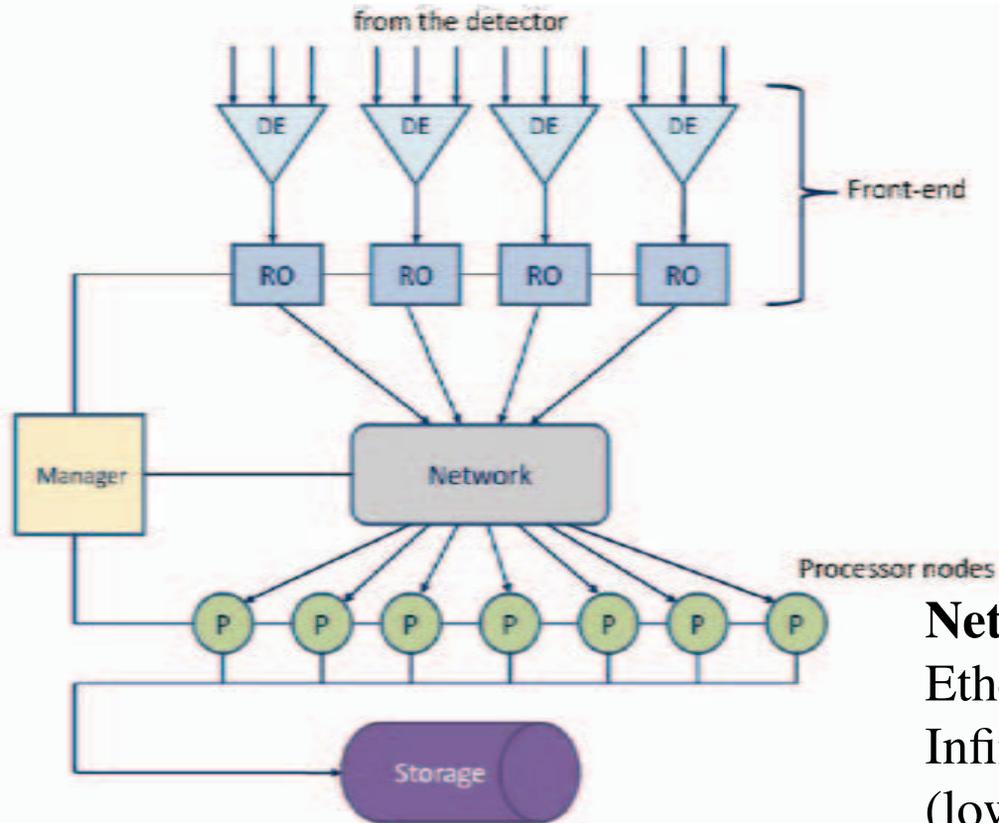


Mu2e physics data set includes selected detector events that contain a candidate high momentum electron ($> 70 \text{ MeV}/c$)

Calibration data sets include min-bias, in situ high-pt positrons, standalone tests, etc.



Mu2e proposal: Streaming DAQ



Network Fabric:
Ethernet: 2011: 100-GBit
InfiniBand: 2012: 100-GBit
(low- latency)

Figure 8.1: A scalable triggerless data acquisition architecture. Detector electronics (DE) sends hit information to readout modules (RO), which transmit data over high-speed optical data links to a network switch that routes data from different parts of the detector to individual processor nodes (P). Events are selected and prepared for storage in the processor nodes. The number of front-end data links and/or the number of processor nodes can be increased to add throughput to the system.



Mu2e DAQ Architecture



Motivations for Streaming Architecture:

- 1) Greater reliance on commodity technology and reconstruction filtering algorithms running on commodity processors “farms”.
- 2) More rapid feedback and development focus on software reconstruction performance.
- 3) Additional efficiency for conversion events. The calorimeter based on-line trigger efficiency is estimated to be 80%. A software trigger component based in tracker information could recover additional conversion events that miss the calorimeter.
- 4) Flexible capacity to collect events with timing (with respect to beam flash) before the nominal start of the live gate, $t = 700\text{nsec}$.
- 5) Flexible and reduced bias data sets to study the tracker, calorimeter, CRV, calibration schemes, and background mechanisms.
- 6) Higher efficiency for positrons in the tracker which opens up additional physics channels with positrons and improves background studies based on positrons.
- 7) Reduces reliance on calorimeter performance for a trigger (more robust against dead channels, blinding, pileup, etc).



ATLAS & CMS Trigger & Readout Structure



≈ 30 Collisions/25ns

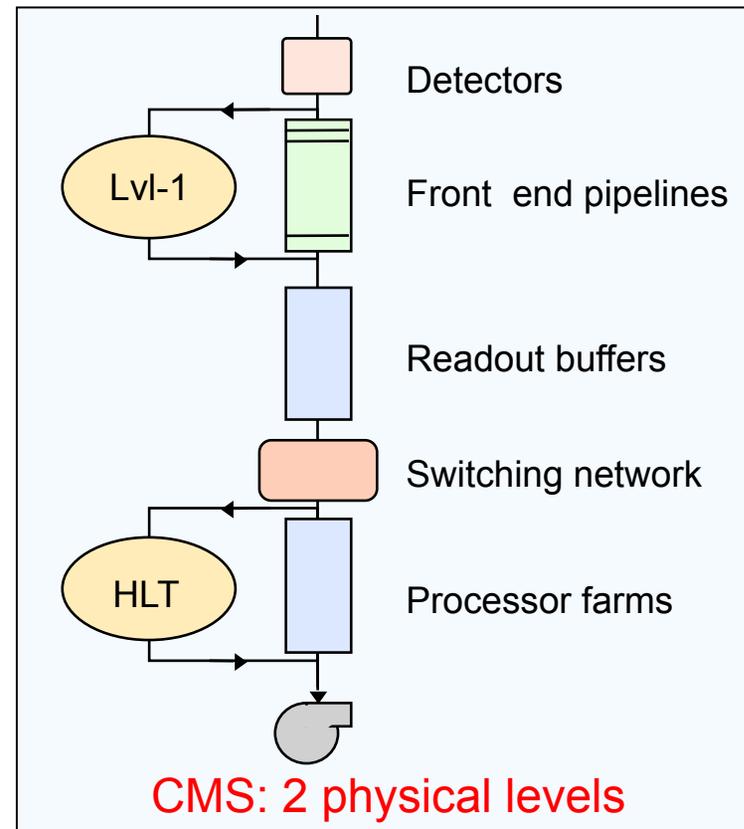
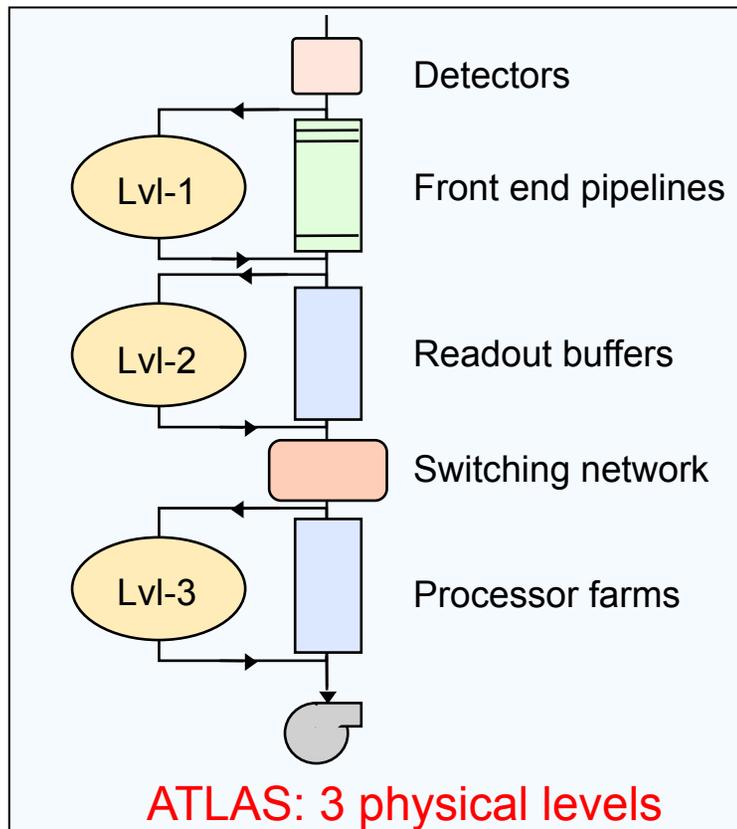
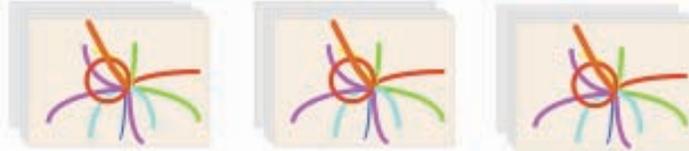
(10^9 event/sec)

10^7 channels

(10^{16} bit/sec)

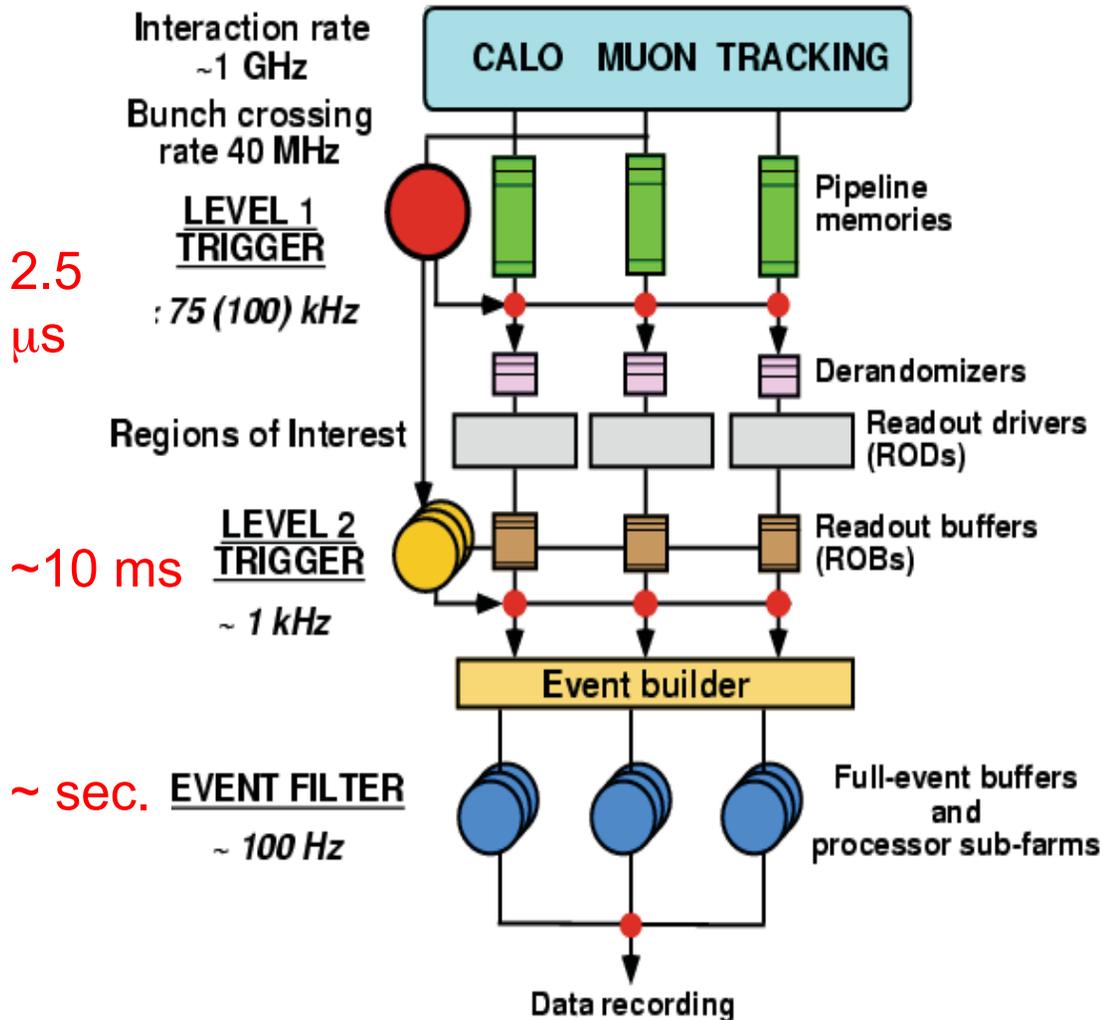
← 25 ns →

Luminosity = 10^{34} cm⁻² sec⁻¹





ATLAS Three Level Trigger Architecture



- **LVL1 decision** made with calorimeter data with coarse granularity and muon trigger chamber data.
 - Buffering on detector
- **LVL2 uses Region of Interest data** (ca. 2%) with full granularity and combines information from all detectors; performs fast rejection.
 - Buffering in ROBs
- **EventFilter** refines the selection, can perform **event reconstruction** at full granularity using latest alignment and calibration data.
 - Buffering in EB & EF



Workshop on Detector R&D

LHC Luminosity Forecast



Year	TeV	OEF	β^*	Nb	lb	ltot	MJ	Peak luminosity	Pile up	pb-1/day	Physics Days	Integrated (fb-1/year)	Total Int (fb-1)	
2010	3.50	0.20	2.00	796	8.0E+10	6.4E+13	36.0	1.886E+32	1.2643	3.3	20.0	0.1	0.07	
2011	3.50	0.25	2.00	796	8.0E+10	6.4E+13	36.0	1.886E+32	1.2643	4.1	240.0	0.98	1.04	
2012	shutdown											0.0	1.0	
2013	6.50	0.20	0.55	796	1.15E+11	9.2E+13	96.1	2.632E+33	17.6429	45.5	180.0	8.2	9.2	
2014	7.00	0.20	0.55	1404	1.15E+11	1.6E+14	182.5	5.000E+33	19.0000	86.4	240.0	20.7	30.0	
2015	7.00	0.20	0.55	2808	1.15E+11	3.2E+14	365.0	1.000E+34	19.0000	172.8	210.0	36.3	66.3	
2016	shutdown											0.0	0.0	66.3
2017	7.00	0.25	0.55	2808	1.15E+11	3.2E+14	365.0	1.000E+34	19.0000	216.0	240.0	51.8	118.1	
2018	7.00	0.28	0.55	2808	1.50E+11	4.2E+14	476.1	1.701E+34	32.3251	411.6	240.0	98.8	216.9	
2019	7.00	0.30	0.55	2808	1.70E+11	4.8E+14	539.6	2.185E+34	41.5198	566.4	210.0	118.9	335.8	
2020	shutdown											0.0	0.0	335.8
2021	7.00	0.20	0.30	2808	1.70E+11	4.8E+14	539.6	4.006E+34	76.1197	692.3	150.0	103.8	439.7	
2022	7.00	0.27	0.25	2808	1.80E+11	5.1E+14	571.3	5.390E+34	102.4060	1257.3	220.0	276.6	716.3	
2023	7.00	0.27	0.25	2808	1.80E+11	5.1E+14	571.3	5.390E+34	102.4060	1257.3	220.0	276.6	992.9	
2024	7.00	0.29	0.25	2808	1.80E+11	5.1E+14	571.3	5.390E+34	102.4060	1350.5	220.0	297.1	1290.0	
2025	7.00	0.29	0.25	2808	1.80E+11	5.1E+14	571.3	5.390E+34	102.4060	1350.5	220.0	297.1	1587.1	
2026	7.00	0.29	0.25	2808	1.80E+11	5.1E+14	571.3	5.390E+34	102.4060	1350.5	220.0	297.1	1884.2	
2027	7.00	0.29	0.25	2808	1.80E+11	5.1E+14	571.3	5.390E+34	102.4060	1350.5	220.0	297.1	2181.3	
2028	7.00	0.29	0.25	2808	1.80E+11	5.1E+14	571.3	5.390E+34	102.4060	1350.5	220.0	297.1	2478.4	
2029	7.00	0.29	0.25	2808	1.80E+11	5.1E+14	571.3	5.390E+34	102.4060	1350.5	220.0	297.1	2775.5	
2030	7.00	0.29	0.25	2808	1.80E+11	5.1E+14	571.3	5.390E+34	102.4060	1350.5	220.0	297.1	3072.6	

Phase 0 → Phase 1 → Phase 2 (high L) →



Requirements for the phases of the upgrades: ~2010-2020



Phase 1:

- This decade will see the initial operation of the LHC and the increase of energy and luminosity towards the design luminosities.
- Goal of extended running in the second half of the decade to collect $\sim 100\text{s/fb}$
- 80% of this luminosity in the last three years of this decade
- About half the luminosity would be delivered at luminosities above the original LHC design luminosity
- Trigger & DAQ systems should be able to operate with a peak luminosity of up to 2×10^{34}



2020-2030 – High Lumi LHC



Phase 2:

- Continued operation of the LHC beyond a few 100/fb will require substantial modification of detector elements
- The goal is to achieve 3000/fb in phase 2
- Need to be able to integrate $\sim 300/\text{fb-yr}$
- Will require new tracking detectors for ATLAS & CMS
- Trigger & DAQ systems should be able to operate with a peak luminosity of up to 5×10^{34} , possibly 10^{35}



CMS Upgrade Trigger Strategy



Constraints

- Output rate at 100 kHz
- Input rate increases x2/x10 (Phase 1/Phase 2) over LHC design (10^{34})
 - Same x2 if crossing freq/2, e.g. 25 ns spacing \rightarrow 50 ns at 10^{34}
- Number of interactions in a crossing (Pileup) goes up by x4/x20
- Thresholds remain ~ same as physics interest does

Example: strategy for Phase 1 Calorimeter Trigger (operating 2016):

- Present L1 algorithms inadequate above 10^{34} or 10^{34} w/ 50 ns spacing
 - Pileup degrades object isolation
- More sophisticated clustering & isolation deal w/more busy events
 - Process with full granularity of calorimeter trigger information
- Should suffice for x2 reduction in rate as shown with initial L1 Trigger studies & CMS HLT studies with L2 algorithms

Potential new handles at L1 needed for x10 (Phase 2: 2020)

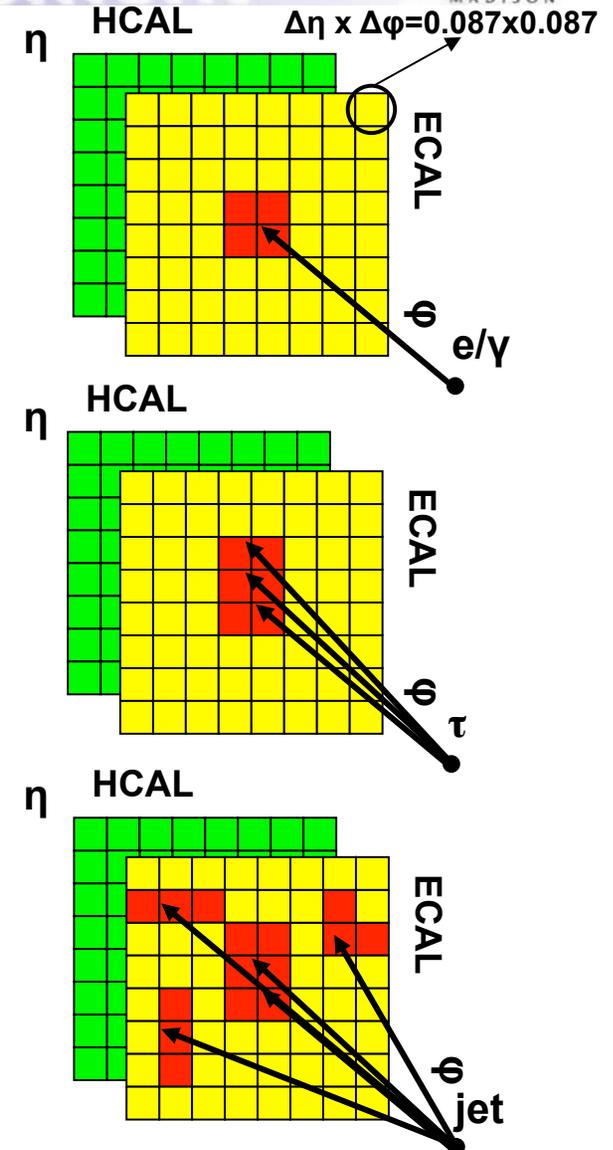
- Tracking to eliminate fakes, use track isolation.
- Vertexing to ensure that the multiple trigger objects come from the same interaction
- Requires finer position resolution for calorimeter trigger objects for matching (provided by use of full granularity cal. trig. info.)



Phase 1 Upgrade Cal. Trigger Algorithm Development



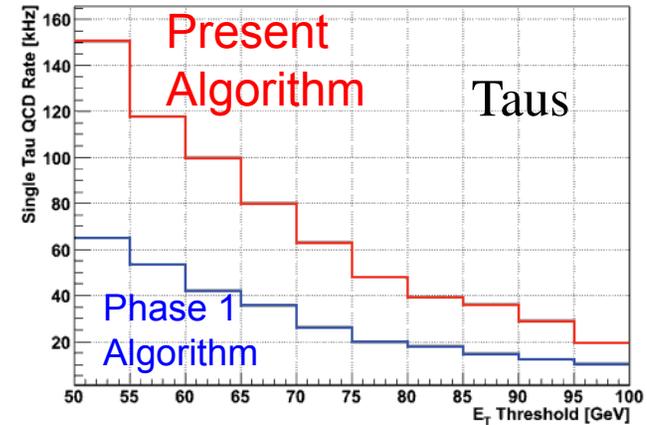
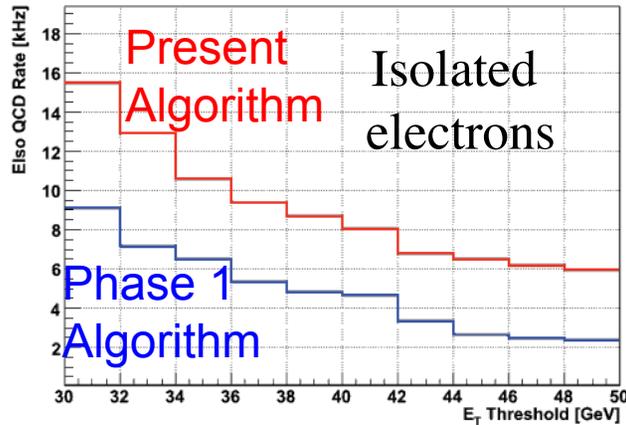
- **Particle Cluster Finder**
 - Applies tower thresholds to Calorimeter
 - Creates overlapped 2x2 clusters
- **Cluster Overlap Filter**
 - Removes overlap between clusters
 - Identifies local maxima
 - Prunes low energy clusters
- **Cluster Isolation and Particle ID**
 - Applied to local maxima
 - Calculates isolation deposits around 2x2,2x3 clusters
 - Identifies particles
- **Jet reconstruction**
 - Applied on filtered clusters
 - Groups clusters to jets
- **Particle Sorter**
 - Sorts particles & outputs the most energetic ones
- **MET,HT,MHT Calculation**
 - Calculates Et Sums, Missing Et from clusters





Upgrade Algorithm Performance: Factor of 2 for Phase I

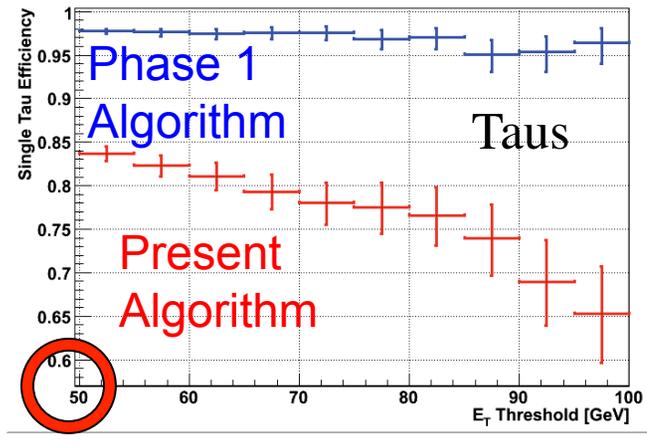
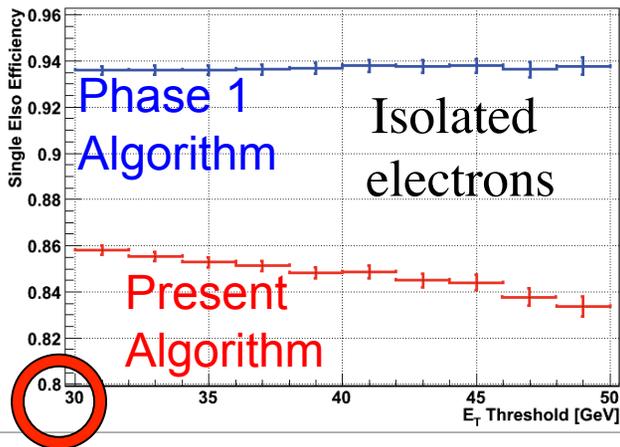
QCD Rate (kHz)



QCD Rate (kHz)

Factor of 2 rate reduction

Efficiency



Efficiency

Higher Efficiency

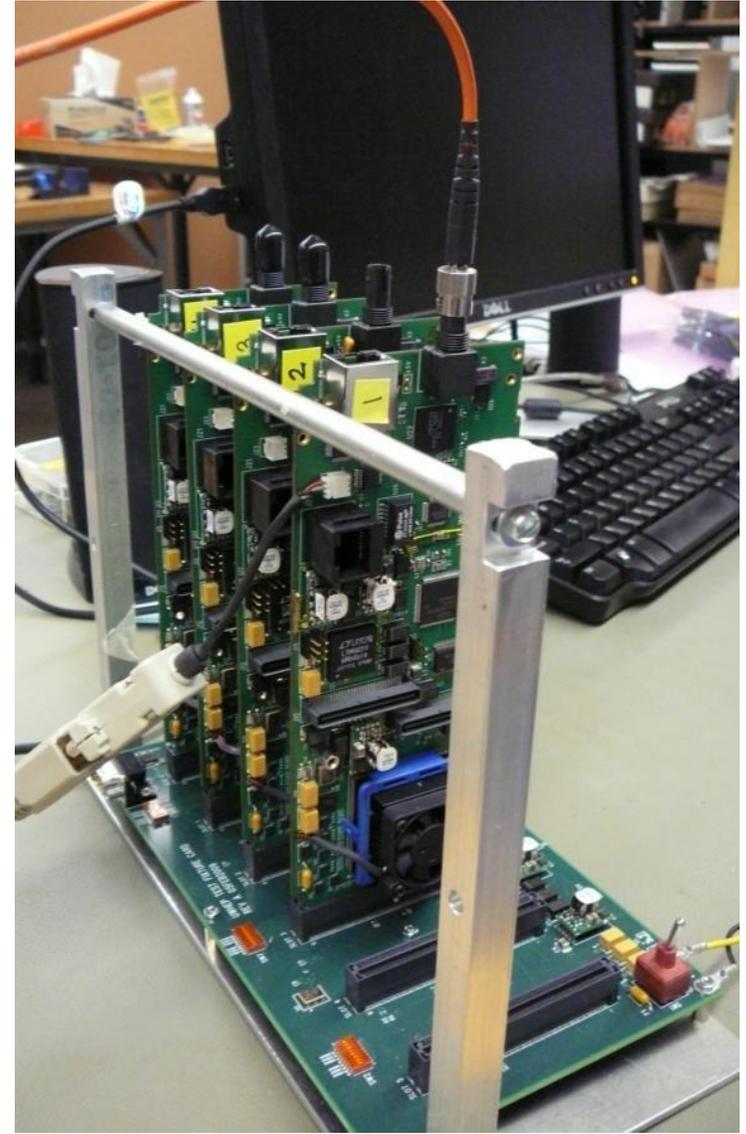
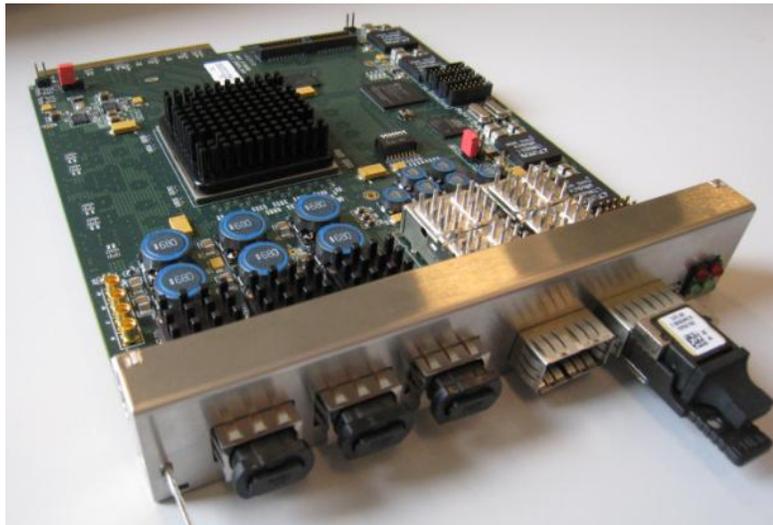


uTCA Calorimeter Trigger Demonstrators



← processing cards with 160 Gb/s input & 100 Gb/s output using 5 Gb/s optical links.

four trigger prototype cards integrated in a backplane fabric to demonstrate running & data exchange of calorimeter trigger algorithms →





CMS CSC Trigger Upgrades



Improve redundancy

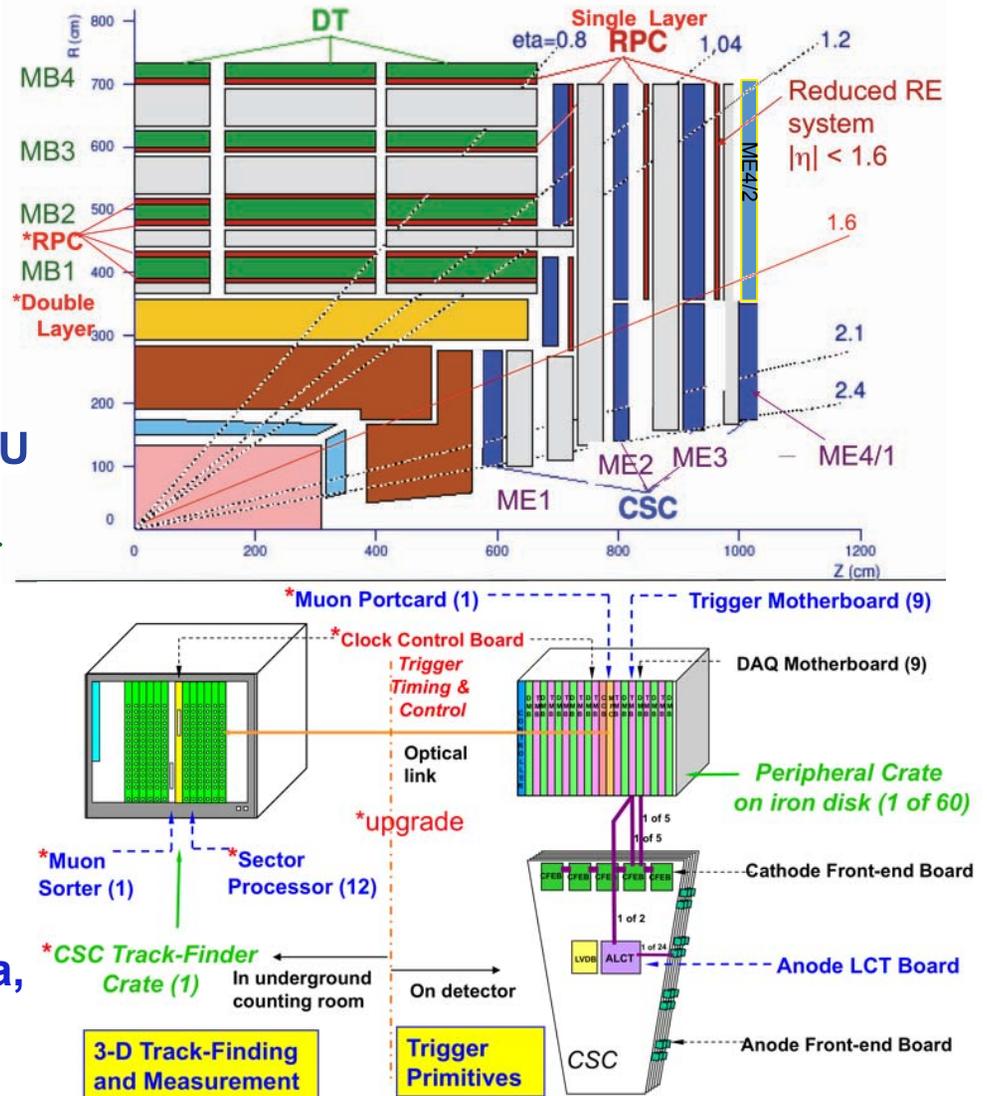
- Add station ME-4/2 covering $\eta=1.1-1.8$
- Critical for momentum resolution

Upgrade electronics to sustain higher rates

- New Front End boards for station ME-1/1
- Forces upgrade of downstream EMU electronics
 - Particularly Trigger & DAQ Mother Boards
- Upgrade Muon Port Card and CSC Track Finder to handle higher stub rate

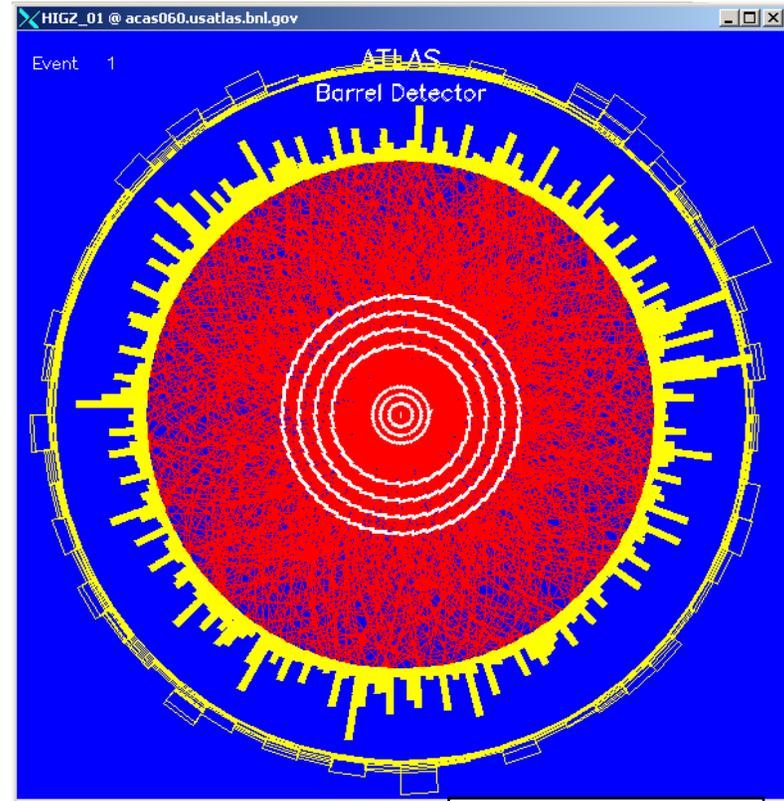
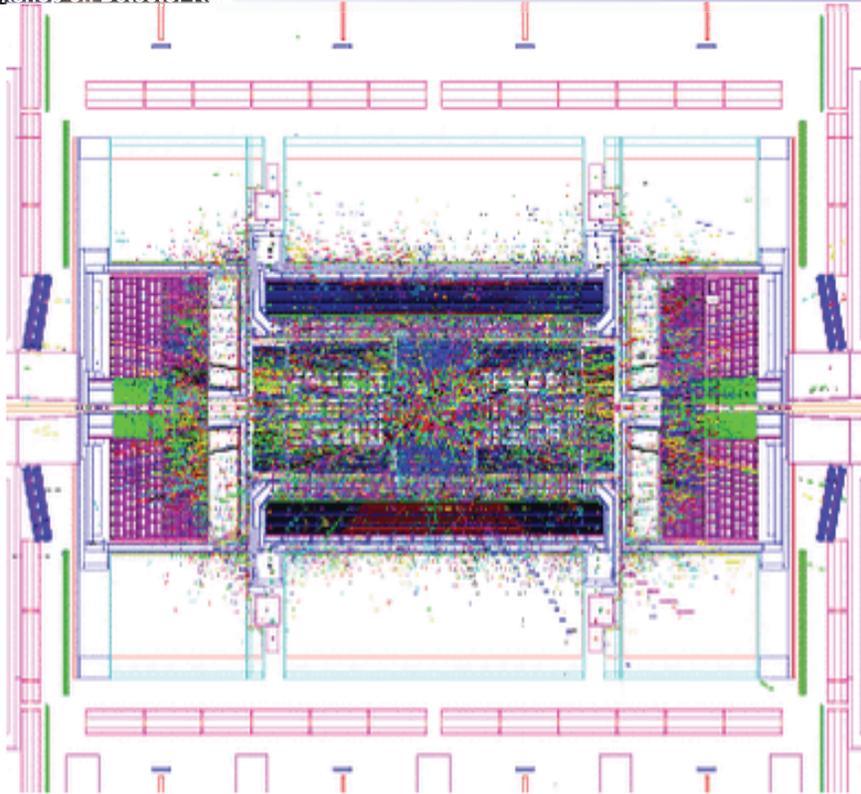
Extend CSC Efficiency into $\eta=2.1-2.4$ region

- Robust operation requires TMB upgrade, unganging strips in ME-1a, new FEBs, upgrade CSCTF+MPC





Expected Pile-up at High Lumi LHC in ATLAS at 10^{35}



$$N_{ch}(|y| \leq 0.5)$$

- 230 min.bias collisions per 25 ns. crossing
- ~ 10000 particles in $|\eta| \leq 3.2$
- mostly low p_T tracks
- requires upgrades to detectors

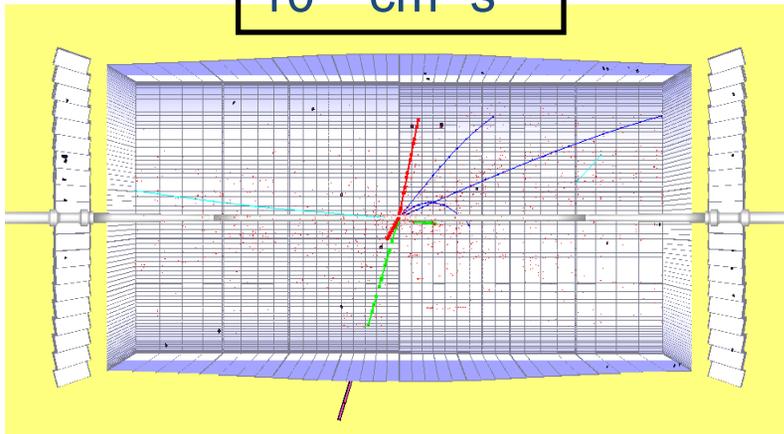


Detector Luminosity Effects

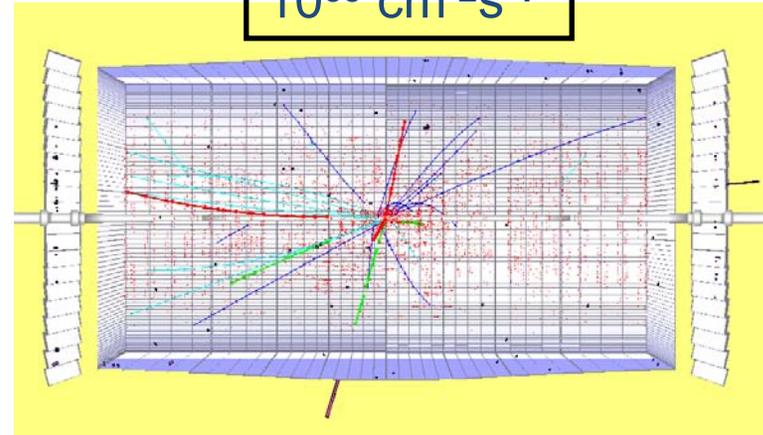


$H \rightarrow ZZ \rightarrow \mu\mu ee$, $M_H = 300$ GeV for different luminosities in CMS

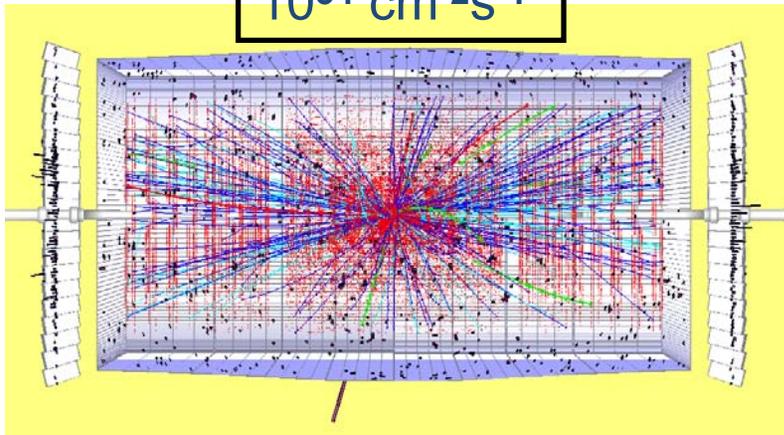
$10^{32} \text{ cm}^{-2}\text{s}^{-1}$



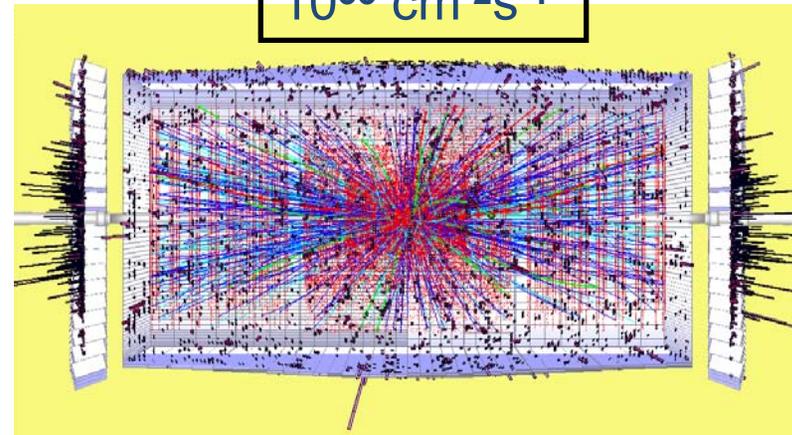
$10^{33} \text{ cm}^{-2}\text{s}^{-1}$



$10^{34} \text{ cm}^{-2}\text{s}^{-1}$



$10^{35} \text{ cm}^{-2}\text{s}^{-1}$





CMS Level-1 Trigger $\rightarrow 10^{35}$



Occupancy

- Degraded performance of algorithms
 - Electrons: reduced rejection at fixed efficiency from isolation
 - Muons: increased background rates from accidental coincidences
- Larger event size to be read out
 - New Tracker: higher channel count & occupancy \rightarrow large factor
 - Reduces the max level-1 rate for fixed bandwidth readout.

Trigger Rates

- Try to hold max L1 rate at 100 kHz by increasing readout bandwidth
 - Avoid rebuilding front end electronics/readouts where possible
 - Limits: $\langle \text{readout time} \rangle (< 10 \mu\text{s})$ and data size (total now 1 MB)
 - Use buffers for increased latency for processing, not post-L1A
 - May need to increase L1 rate even with all improvements
 - Greater burden on DAQ
- Implies raising E_T thresholds on electrons, photons, muons, jets and use of multi-object triggers, unless we have new information \Rightarrow Tracker at L1
 - Need to compensate for larger interaction rate & degradation in algorithm performance due to occupancy



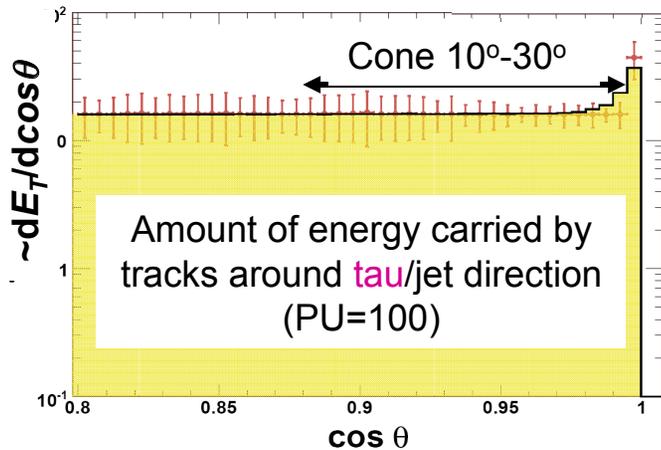
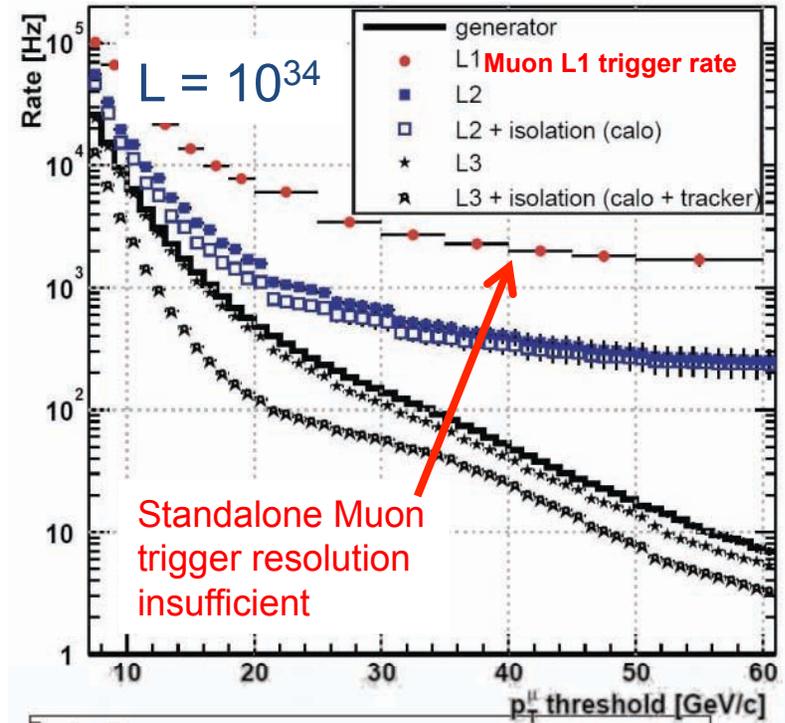
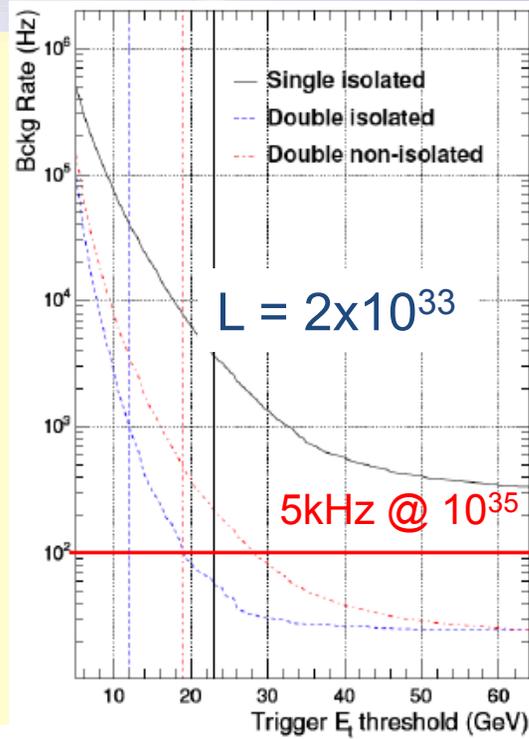
Workshop on Detector R&D

Tracking needed for L1 trigger

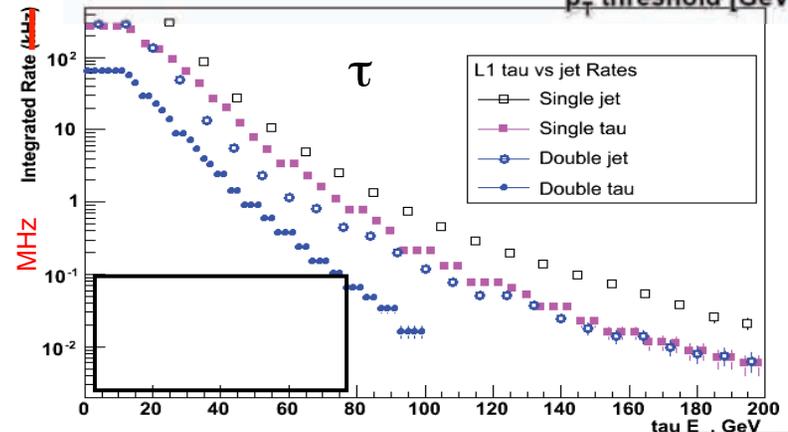


Single electron trigger rate

Isolation criteria are insufficient to reduce rate at $L = 10^{35} \text{ cm}^{-2} \cdot \text{s}^{-1}$



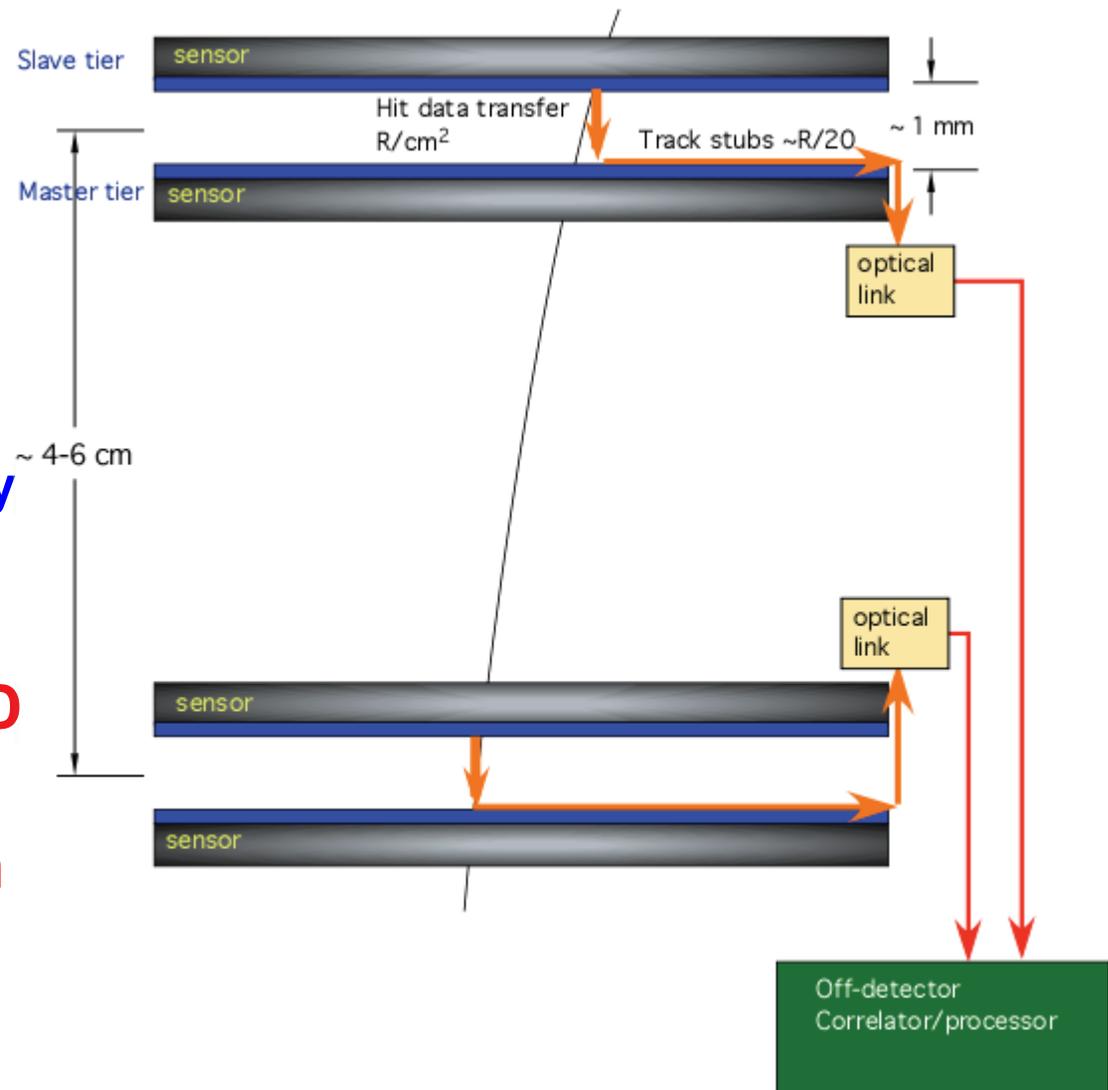
We need to get another x200 (x20) reduction for single (double) tau rate!





The Track Trigger Problem

- Need to gather information from 10^8 pixels in 200m^2 of silicon at 40 MHz
- Power & bandwidth to send all data off-detector is prohibitive
 - Local filtering necessary
 - Smart pixels needed to locally correlate hit P_t information
- Studying the use of 3D electronics to provide ability to locally correlate hits between two closely spaced layers





3D Interconnection

Key to design is ability of a single IC to connect to both top & bottom sensor

Enabled by “vertical interconnected” (3D) technology

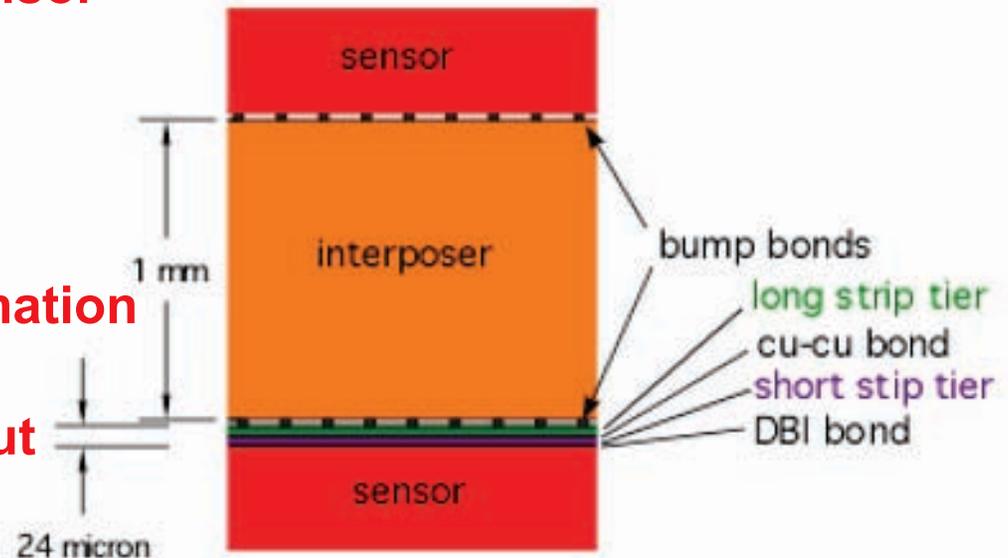
A single chip on bottom tier can connect to both top and bottom sensors – locally correlate information

Analog information from top sensor is passed to ROIC (readout IC) through interposer

One layer of chips

No “horizontal” data transfer necessary – lower noise and power

Fine Z information is not necessary on top sensor – long (~1 cm vs ~1-2 mm) strips can be used to minimize via density in interposer





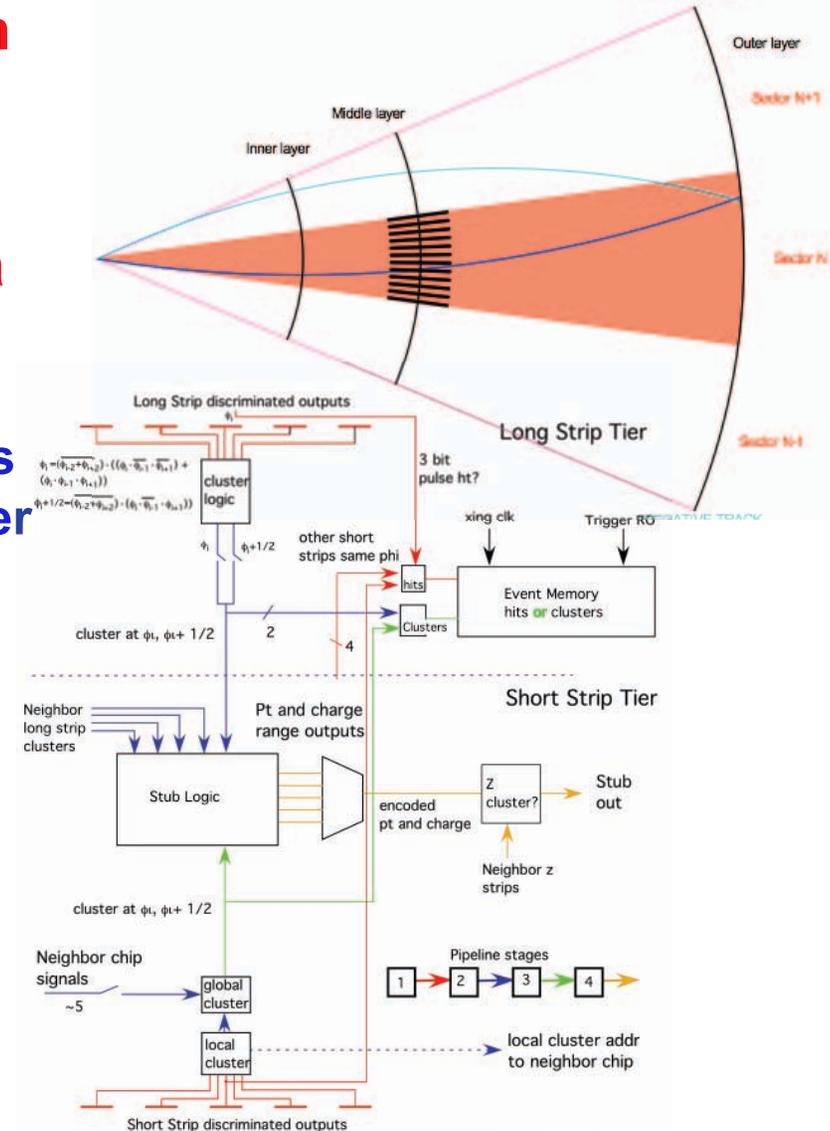
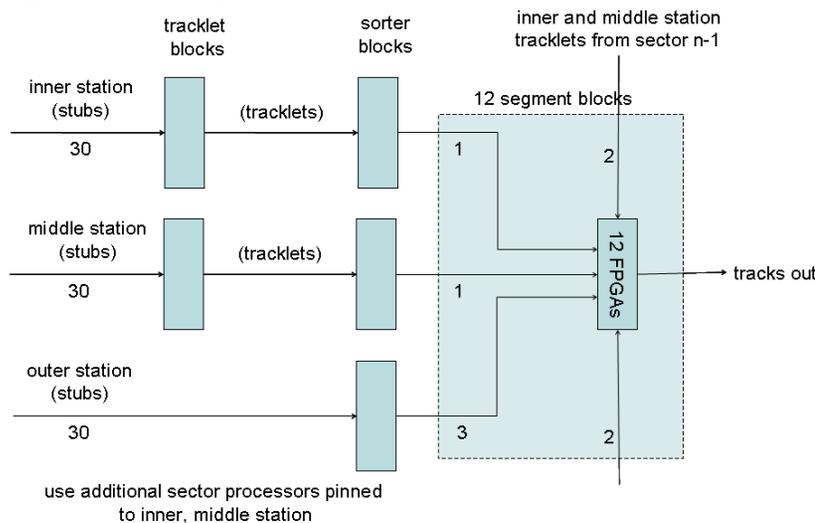
Track Trigger Architecture

Readout designed to send all hits with $P_t > \sim 2$ GeV to trigger processor

High throughput – micropipeline architecture

Readout mixes trigger and event data
Tracker organized into phi segments

- Limited FPGA interconnections
- Robust against loss of single layer hits
- Boundaries depend on p_t cuts & tracker geometry



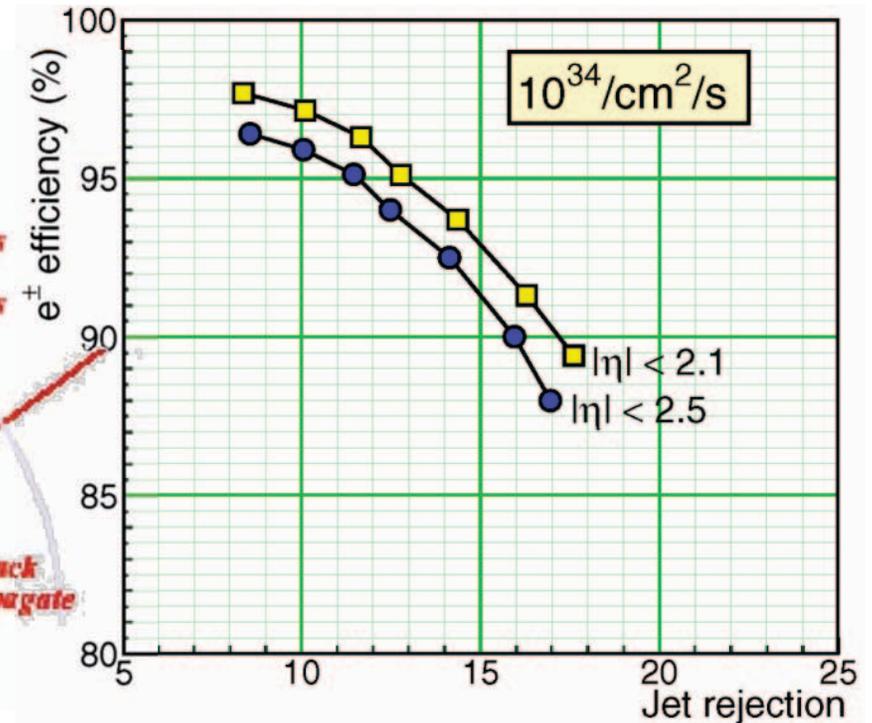
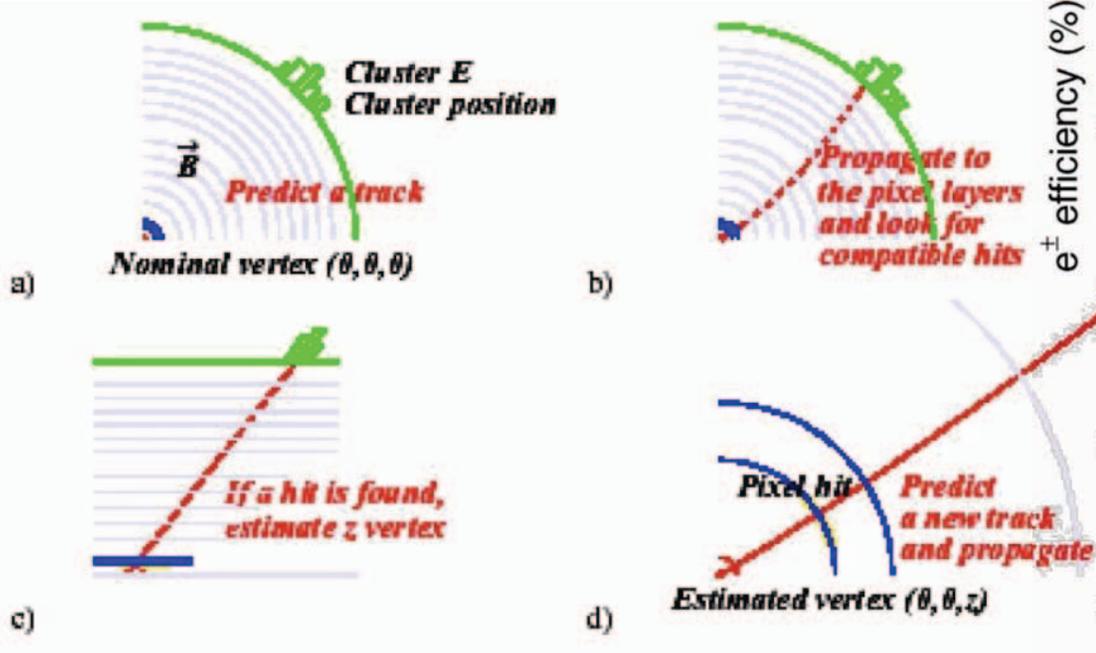


Tracking for electron trigger



Present CMS electron HLT

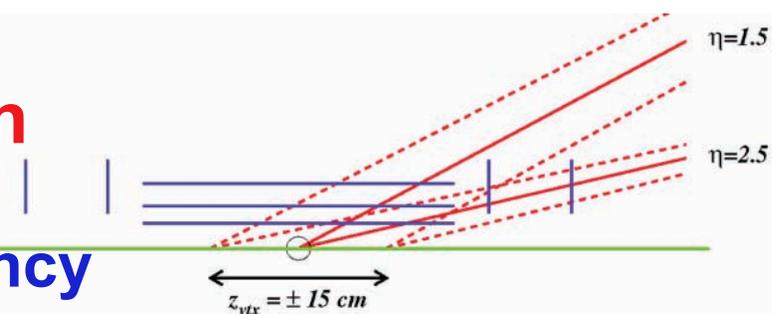
- C. Foudas & C. Seez



Factor of 10 rate reduction

γ : only tracker handle: isolation

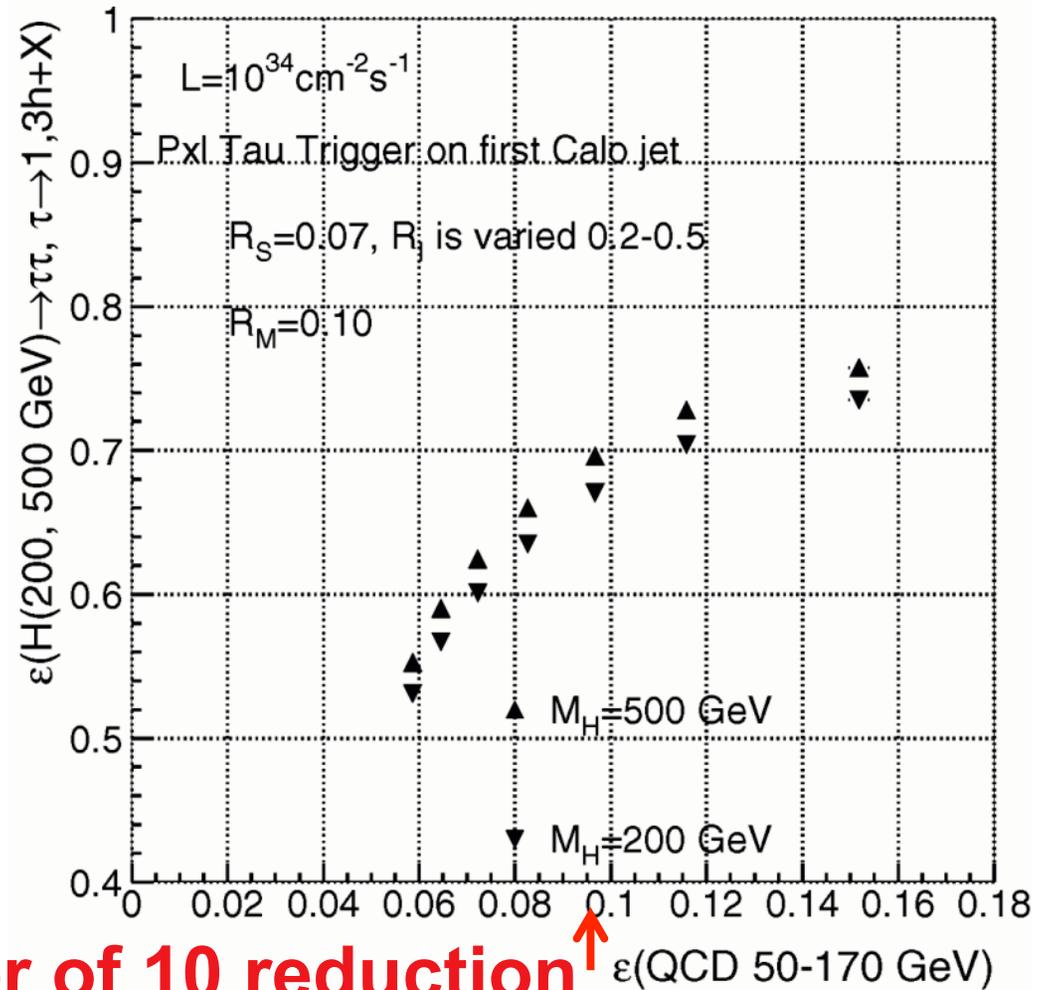
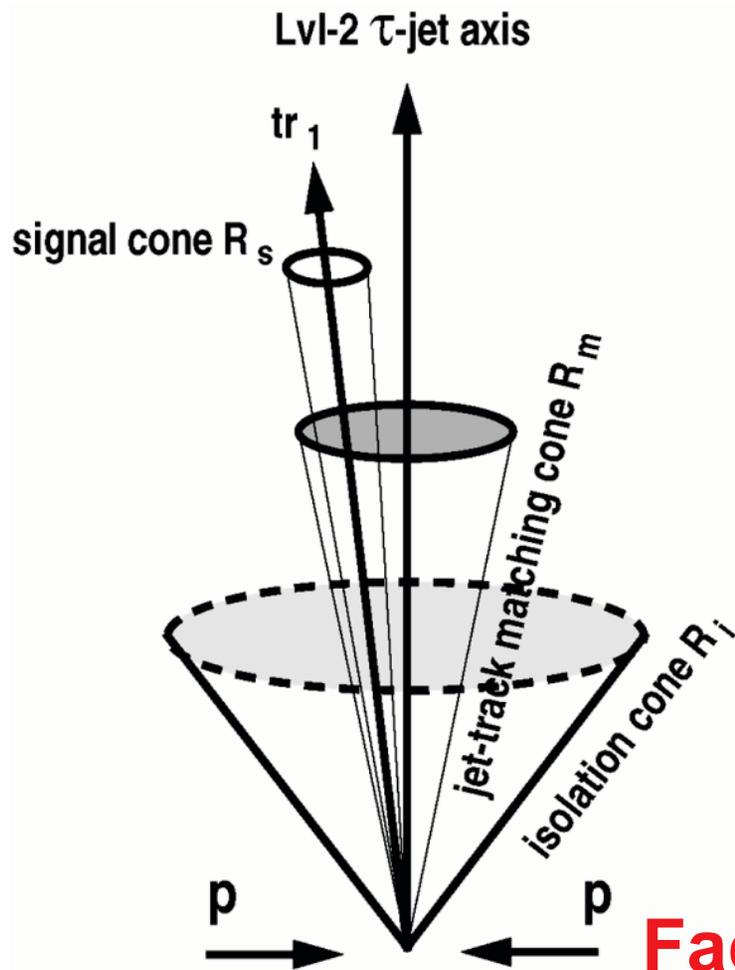
- **Need knowledge of vertex location to avoid loss of efficiency**





Tracking for τ -jet isolation

τ -lepton trigger: isolation from pixel tracks outside signal cone & inside isolation cone





CMS L1 Track Trigger for Muons



Combine with L1 μ trigger as is now done at HLT:

- Attach tracker hits to improve P_T assignment precision from 15% standalone muon measurement to 1.5% with the tracker
 - Improves sign determination & provides vertex constraints
- Find pixel tracks within cone around muon track and compute sum P_T as an isolation criterion
 - Less sensitive to pile-up than calorimetric information *if* primary vertex of hard-scattering can be determined (~100 vertices total at SLHC!)

To do this requires η - ϕ information on muons finer than the current 0.05 - 2.5°

- No problem, since both are already available at 0.0125 and 0.015°

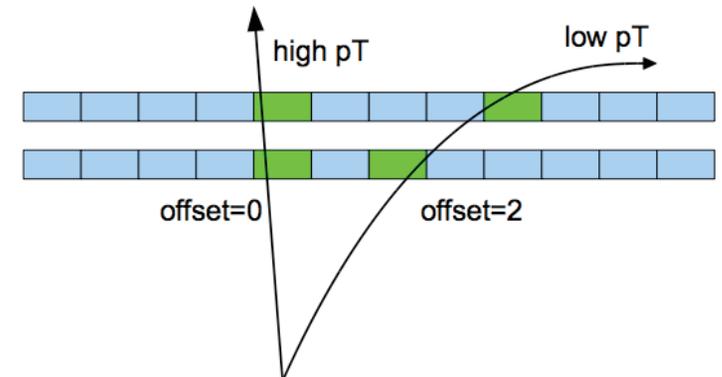


ATLAS Trigger Upgrades



Various projects being pursued:

- **Track trigger**
 - Fast Track Finder (FTK), hardware track finder for ATLAS (at L1.5)
 - ROI based track trigger at L1
 - Self seeded track trigger at L1
- **Combining trigger objects at L1 & topological "analysis"**
- **Full granularity readout of calorimeter**
 - requires new electronics
- **Changes in muon systems (small wheels), studies of an MDT based trigger & changes in electronics**
- **Upgrades of HLT farms**



Some of the changes are linked to possibilities that open when electronics changes are made (increased granularity, improved resolution & increased latency)



ATLAS Cal. Trigger upgrade



Phase I: upgrade current L1Calo

- FPGA-based MCM replacement for PreProcessor (?)
- Augment EM/Had and Jet/Energy processors with CMM++ to add topological algorithm capabilities
 - Replacement for present trigger data Common Merger Module

Phase II: Replace L1Calo with 2-level system

- Full digital readout of LAr, Tile data to Readout Drivers (RODs) in underground counting room (USA15)
- "Level 0": Synchronous, fixed latency, Topological algorithms with calorimeters + muon ROIs
 - Uses trigger towers (0.1×0.1) w/finer $\eta \times \phi$, depth segmentation.
- "Level 1": Asynchronous, longer latency, access to full resolution calorimeter data, Topological algorithms with calo, muon and ID ROIs
 - Improved ID of isolated electrons, hadrons identified by L0
 - Aim for similar performance to present L2



ATLAS Muon Trigger Upgrade



MDT precision can be used for L1 sharpening

- Present ATLAS muon trigger based on RPCs only.
- Use RPC L1 trigger as “seed”. MDTs only verify p_T on request from RPC
 - No stand-alone trigger of Monitored Drift Tubes
- Use RPC hits to define a search road for corresponding MDT hits

Need extra latency of $\sim 2 \mu\text{s}$ (Phase 2)

Benefits:

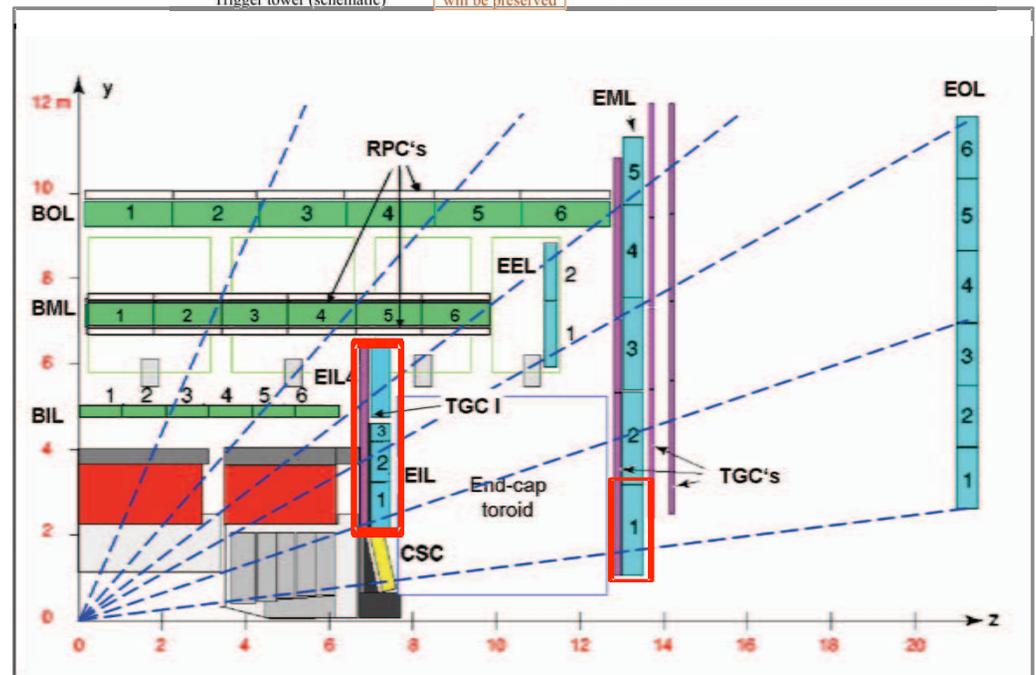
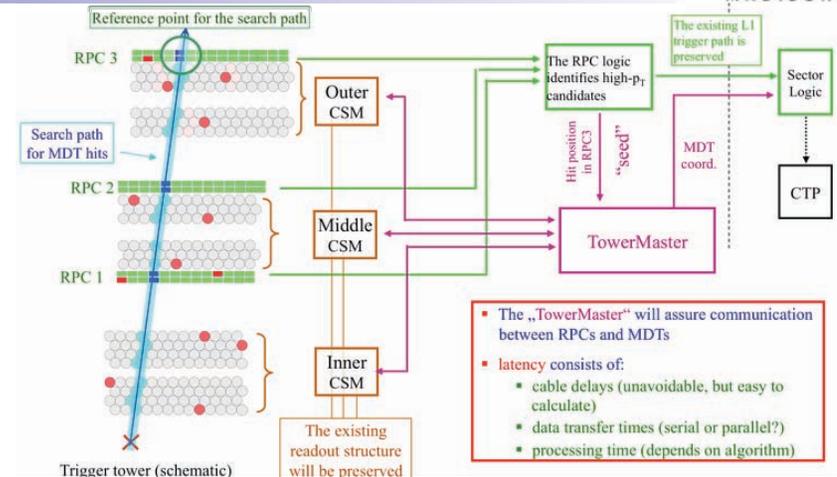
- No additional trigger chambers required in Barrel
- No interference with normal readout

Hardware consequences: concept needs

- rebuilding of MDT electronics
- modification of parts of RPC electronics (PADs, Sector Logic).

Requires new chips & boards:

- New front end board (mezzanine)
- New Chamber Service Module
- New architecture of RPC/TowerMaster
- interface to RPC readout





ATLAS FastTracker (FTK)



For Phase 1:

Dedicated hardware processor completes GLOBAL track reconstruction by beginning of level-2 processing.

- Allows very rapid rejection of most background, which dominates the level-1 trigger rate.
- Frees up level-2 farm to carry out needed sophisticated event selection algorithms.

Addresses two time-consuming stages in tracking

- **Pattern recognition** – find track candidates with enough Si hits
 - 10^9 prestored patterns simultaneously see each silicon hit leaving the detector at full speed.
- **Track fitting** – precise helix parameter & χ^2 determination
 - Equations linear in local hit coordinates give near offline resolution

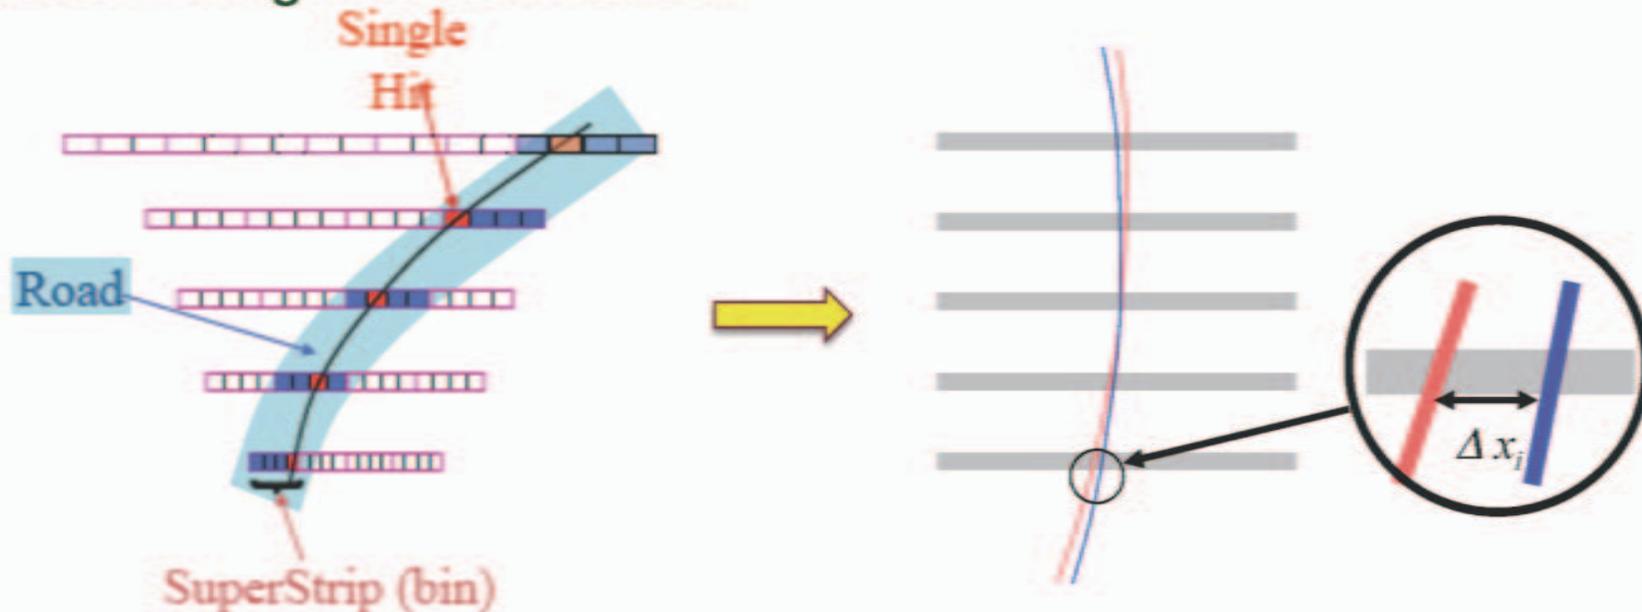


Workshop on Detector R&D

ATLAS FTK Approach



Use hardware to perform the global tracking in two steps
pattern recognition and track fit



Pattern recognition in coarse resolution
(superstrip → road)

Track fit in full resolution (hits in a road)
 $F(x_1, x_2, x_3, \dots) \sim a_0 + a_1 \Delta x_1 + a_2 \Delta x_2 + a_3 \Delta x_3 + \dots = 0$

**Design: FTK completes global tracking in 25 μ sec at 3×10^{34} .
Current level-2 takes 25 msec per jet or lepton at 3×10^{34} .**



ATLAS L1 Track Trigger Design Options for Phase 2



Region Of Interest based Track Trigger at L1

- uses ROIs from L1Calo & L1Muon to seed track finding
- has a large impact on the Trigger architecture
 - requires significantly lengthened L1 pipelines and fast access to L1Calo and L1Muon ROI information
 - could also consider seeding this with an early ("Level-0") trigger, or sending a late ("Level-1.5") track trigger
- smaller impact on Silicon readout electronics

Self-Seeded Track Trigger at L1

- independent of other trigger information
- has a large impact on Silicon readout electronics
 - requires fast access to Silicon detector data at 40 MHz
- smaller impact on the Trigger architecture



ATLAS RoI Tracking Trigger



L0 similar to current L1-Calor & L1-Muon defines regions of interest (Rols)

- There is no inner detector (tracking) information in the RoI definition

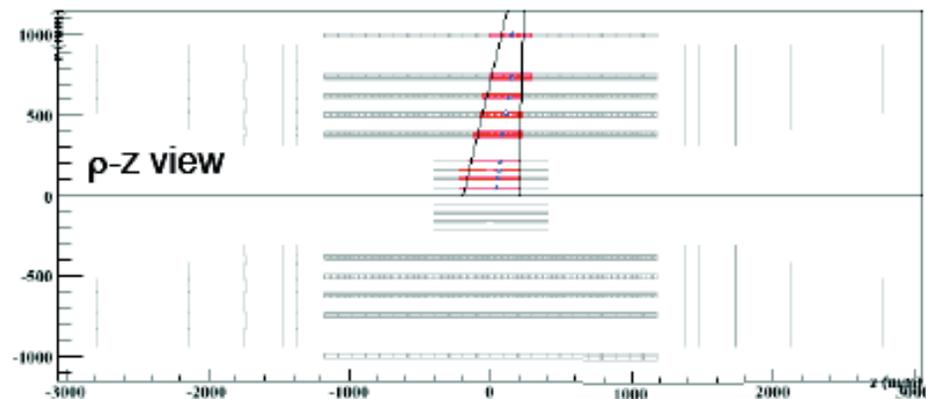
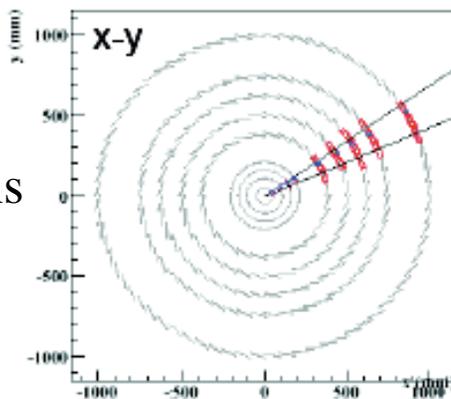
RoI defines an eta-phi region for strips & pixel information to be extracted

L1 uses inner detector information from Rols that were defined in L0

- Can also do a detailed correlation with outer detector

RoI: $\Delta\phi=0.2$, $\Delta\eta=0.2$ at Calo $\Delta z=40\text{cm}$ at beam line

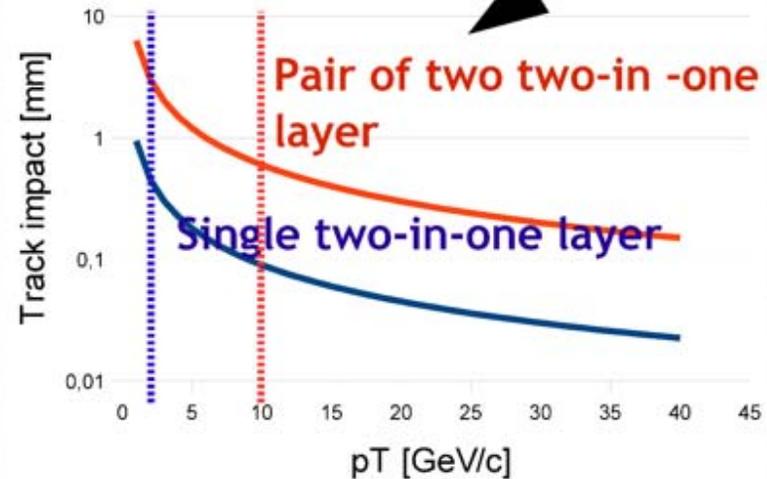
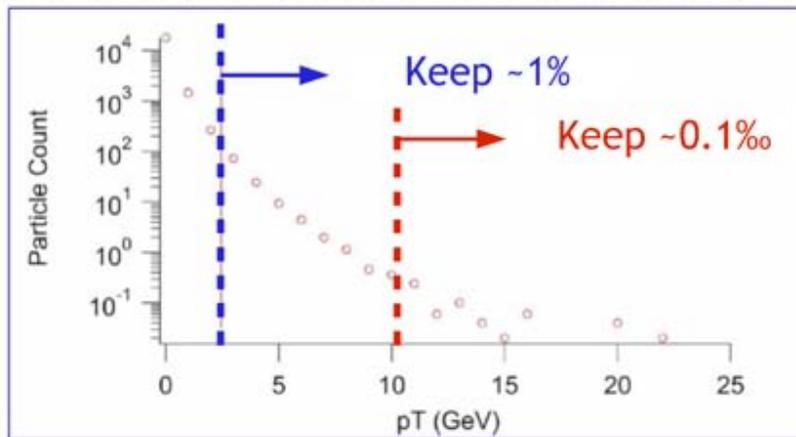
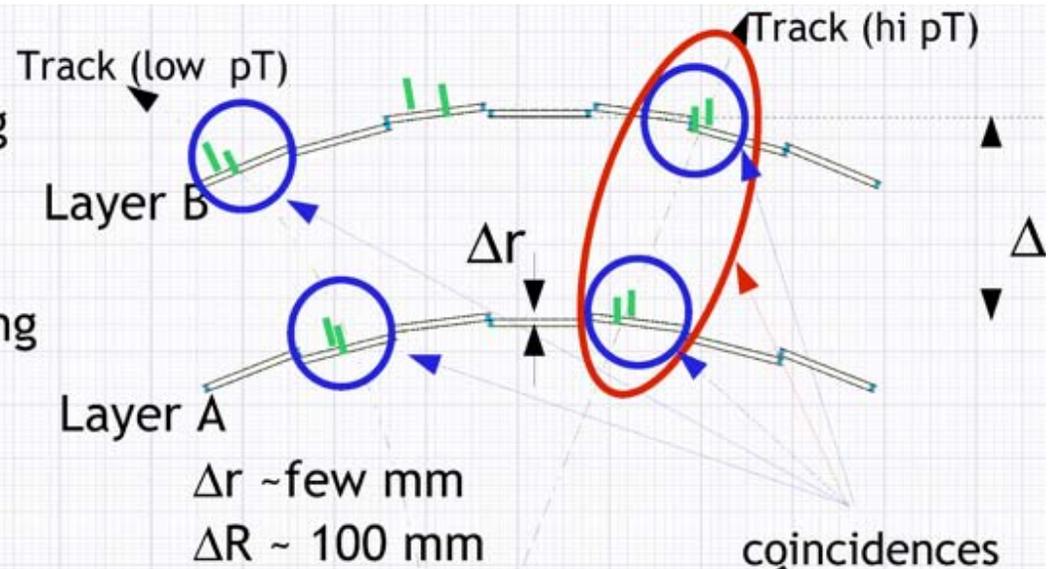
Example of RoI: Contains $\approx 1\%$ of tracker modules





ATLAS Self-Seeded L1 Track Trigger with Doublet Layers

- Moderate pT dependent discrimination of hits using coincidences in closely spaced double layers
- High pT discrimination using coincidences between several doublet layers
- Has to operate at full BCO frequency (40 MHz)

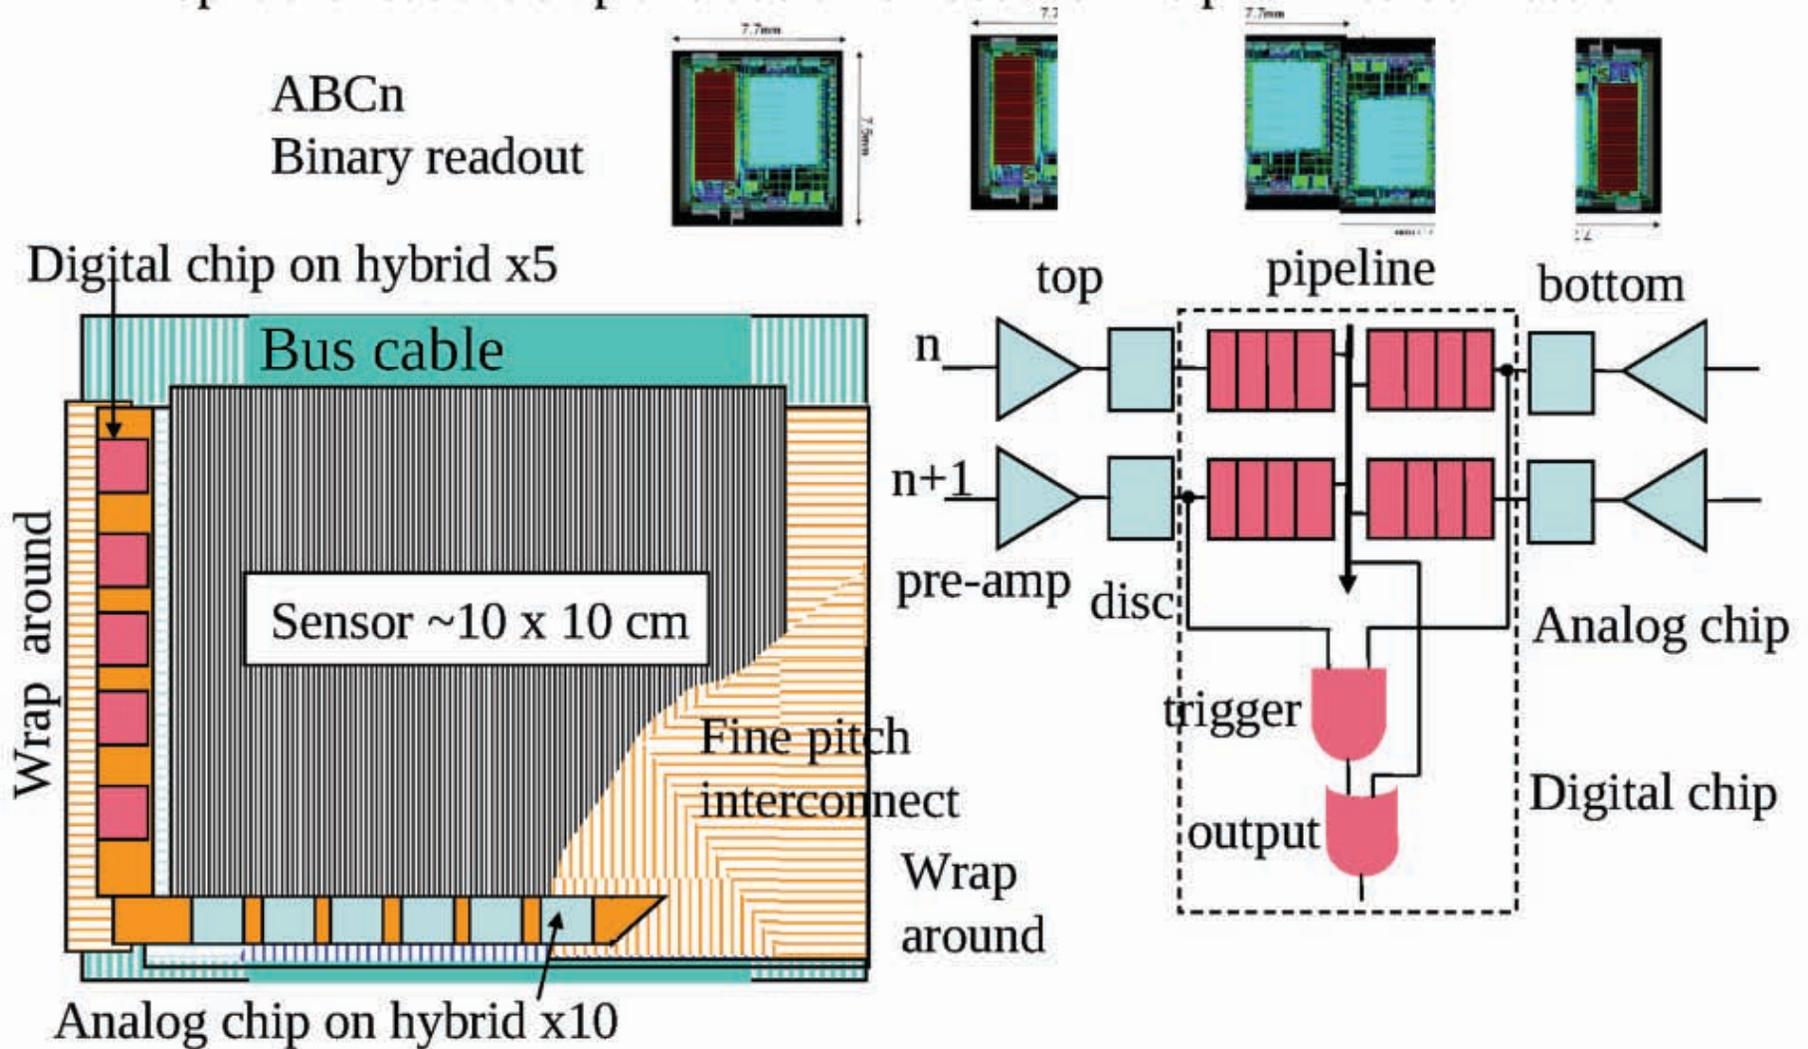




ATLAS Self-Seeded L1 Track Trigger: One possible solution



Split the readout chip and add an embedded fine pitch interconnection





Upgrade DAQ



Phase 2 Network bandwidth at least 5-10 times LHC

- Assuming L1 trigger rate same as LHC
- Increased Occupancy
- Decreased channel granularity (esp. tracker)

CMS DAQ Component upgrades

- Readout Links: replace existing SLINK (400 MB/s) with 10 Gbit/s
- Present custom Front End Detector Builder & Readout Unit Builder replaced with modern network technology & multi-gigabit link network switch
- Higher Level Trigger CPU Filter Farm estimates:
 - 2010 Farm = 720 Dual Quad Core Xeon 5430 (2.66 GHz)
 - 2012 Farm = 3X present farm
 - 2016 Farm = 3X 2012 farm
 - Requires upgrades to network infrastructure



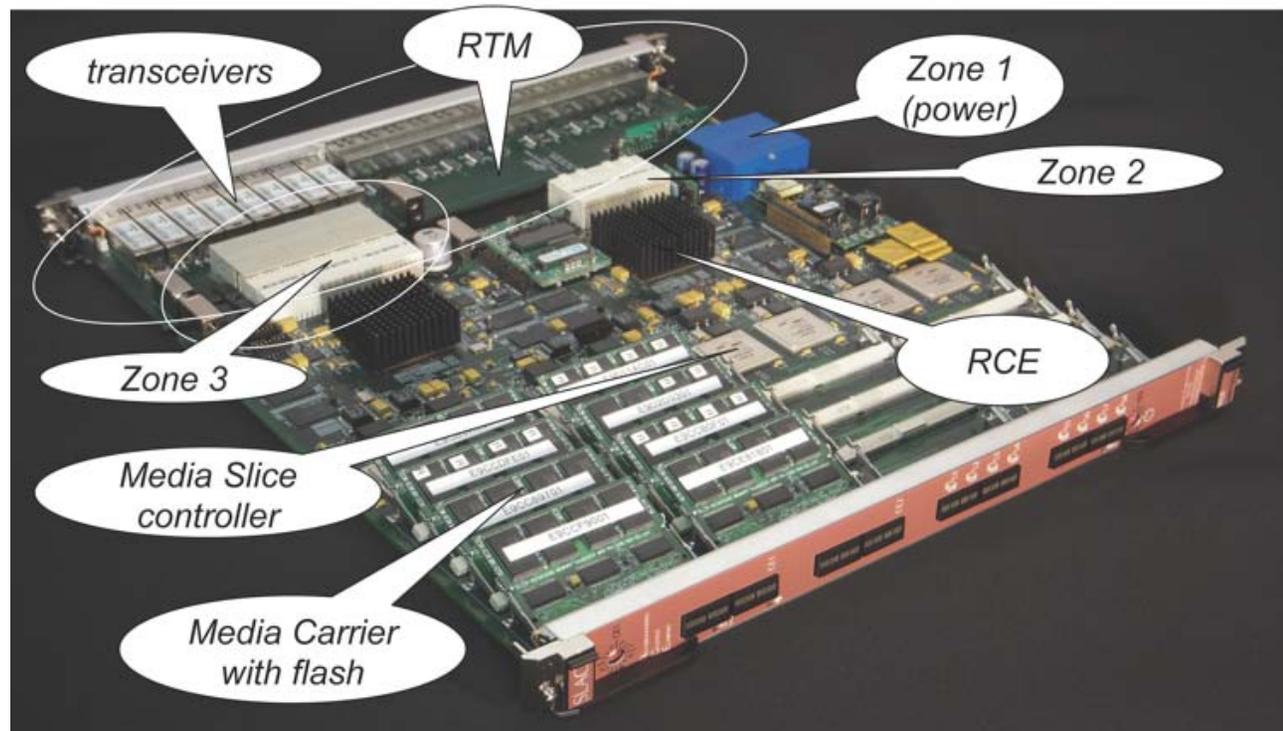
ATLAS Upgrade DAQ



One project explores full capabilities of large modern FPGAs for versatile generic DAQ with its core effort named as Reconfigurable Cluster Element (RCEs), implemented on ATCA platform.

First generation boards in use on SLAC LCLS experiments, LSST DAQ, PetaCache proj. Studying possible use for ATLAS pixel upgrade

Board shown here with 1TB FlashRAM for PetaCache project





Workshop on Detector R&D

High Rate Experiment R&D



Very significant challenges to operate trigger & DAQ systems for high rate experiments, particularly examples shown for ATLAS & CMS

Very substantial assets to bring to bear on these challenges from commercial world: μ TCA, FPGAs, high speed links (transceivers).

Exploiting these assets enables physics input to drive much more precise selection of events and processing of a much higher volume of data.

- e.g. a level-1 tracking trigger**

There is considerable technical difficulty involved in successfully exploiting these advances in technology and implementing them in running experiments in a controlled and adiabatic manner.



Acknowledgements



Thanks to the following for slides and information:

- **Darin Acosta, Bob Bernstein, Sridhara Dasu, Su Dong, Hal Evans, Costas Foudas, Tom Gorski, Reiner Hauser, Eric Hazen, Greg Iles, Andy Lankford, Jim Linneman, Eliot Lipeles, Ron Lipton, Jordan Nash, Abe Seiden, Mel Shochet, Sam Silverstein, Bob Tschirhart**